EVIDENCE AND EPISTEMIC NORMATIVITY

Daniel Buckley

Submitted to the faculty of the University Graduate School

in partial fulfillment of the requirements

for the degree

Doctor of Philosophy

in the Department of Philosophy,

Indiana University

August, 2022

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Doctoral Committee

_____

Adam Leite, Ph.D.

_____

Kate Abramson, Ph.D.

_____

Jordi Cat, Ph.D.

_____

Gary Ebbs, Ph.D.

July 25, 2022

For my parents, John and Jean Buckley

Acknowledgements

This dissertation would not have been completed were it not for the support and guidance of many people. First and foremost, I would like to thank my chair, Adam Leite. Adam's insightful feedback helped shape and improve this project in countless ways. I am extremely grateful for his support and encouragement over the years, and I have enjoyed trying to become a better philosopher under his direction. I would also like to thank the rest of my committee: Kate Abramson, Gary Ebbs, and Jordi Cat. Katy and Gary have always been there to discuss philosophy with me, and I've benefitted enormously from their suggestions and feedback. Jordi was also a valuable resource as I was completing my minor requirement. Finally, while not an official member of my committee, Frederick Schmitt has provided me with invaluable feedback and advice.

I don't think I would have made it through graduate school were it not for my friends and fellow grad students in the Philosophy department at IU. I'll look back fondly on my years in Bloomington in large part because of the friendship of an amazing group of people. I'd especially like to thank: Levi Tenen, Sean Murphy, Ivan Verano, Uri Eran, Andrew Smith, Dave Fisher, Daniel Linsenbardt, Kevin Mills, Wade Munroe, Alex and Kristin Buchinski, Sarah Adams, and Emmalon Davis. Thanks to you all for making life in Bloomington so wonderful.

I would like to thank my parents, John and Jean Buckley, to whom this dissertation is dedicated. Their love and support for me has been constant, regardless of the paths that I have chosen to walk down, and for that I am forever grateful. Last but certainly not least, I would like to thank Sarah Dunevant. Sarah has been by my side during some of the hardest moments in graduate school. Her love and companionship got me through those moments. I simply wouldn't know what to do without her. Thank you from the bottom of my heart.

Daniel Buckley

EVIDENCE AND EPISTEMIC NORMATIVITY

Abstract: Evidence is often taken to be "normative" for doxastic agents. For instance, we are told to apportion our beliefs to the evidence, to not believe a claim without seeking out countervailing evidence, and so on. But what accounts for the normativity of evidence? This dissertation is devoted to answering this question. In order to answer it, I develop a novel approach to the theory of epistemic normativity. According to my approach, epistemic norms structure a social practice of *epistemic accountability*. This approach shares affinities with Strawsonian attempts to elucidate moral responsibility by considering the conditions under which it's appropriate to subject a person to the "reactive attitudes" (e.g. resentment and indignation). However, when it comes to *epistemic* (as opposed to *moral*) accountability, I argue that the relevant attitudinal responses need not involve reactive emotions. Moreover, what I seek to elucidate by appealing to these attitudinal responses is not responsible agency, but rather the content and normative significance of epistemic norms.

Crucial to my approach is a *distinctly epistemic* way of holding a person accountable for their doxastic attitudes. To hold a person epistemically accountable, on my approach, is to modify trust in the person in a particular way. For instance, someone might cease to take a person's words at face value when it comes to a certain topic, or someone might be unwilling to rely on another as a testimonial source of information. I argue that our practice of epistemic accountability is a legitimate social practice. I then go on to consider what norms of belief structure its activities. I argue that a number of norms of belief are implicit in our practice of epistemic accountability, including *evidential* norms and *knowledge* norms. I ultimately argue that our acceptance of epistemic norms is itself grounded in the fact that we participate in a social practice wherein we hold each other epistemically accountable.

Contents

Introduction

1. Epistemic Normativity: An Introduction to the Problem

The problem of epistemic normativity is an instance of the general problem of normativity. As it turns out, it takes a fair amount of philosophical work to pin down what the "problem" is supposed to be, exactly. However, Christine Korsgaard gives us a good place to start. In *The Sources of Normativity,* Korsgaard was concerned specifically with the demands of *morality,* and in particular its ability to generate overriding obligations. What prompted her to investigate this was a certain question that she called the "normative question". Imagine an account of the claims of morality which reveals that ethical standards were invented by politicians to keep us in line, or that moral standards have some kind of biological basis. What would be the problem with such accounts? According to Korsgaard,

> The answer lies in the fact that ethical standards are *normative*. They do not merely *describe* a way in which we in fact regulate our conduct. They make *claims* on us; they command, oblige, recommend, or guide. Or at least, when we invoke them ,we make claims on one another. When I say that an action is right I am saying that you ought to *do* it; when I say something is good I am recommending it as worthy of your choice. The same is true of the other concepts for which we seek a philosophical foundation. Concepts like beauty, knowledge, and meaning, as well as virtue and justice, all have a normative dimension, for they tell us what to think, what to like, what to say, what to do, and what to be. And it is the force of these normative claims – the right of these concepts to give laws to us – that we want to understand…When we seek a philosophical foundation for morality we are not looking merely for an explanation of moral practices. We are asking what *justifies* the claims that morality makes on us.[1]

The dissatisfaction that Korsgaard is registering with the aforementioned accounts might be difficult to appreciate at first. But place yourself in the shoes of someone who is facing a difficult moral decision. Say that you know that some course of action would be the right thing to do, but

---

[1] Korsgaard (1996) pp. 8-10

it would be *so much easier* to just do something else, and no one will know the difference anyway! Why do the right thing? If morality is just an invention of the powerful, or if it's just nature's way of promoting the survival of our species, then the authority of the initial claim may begin to appear fraudulent, or at least attenuated. One might then start to wonder, perhaps in a Nietzschean spirit, whether we've internalized certain psychological tendencies associated with this thing we call "morality" which are downright unhealthy or inhibiting.

This is what prompts Korsgaard's attempt to vindicate the claims of morality. In this dissertation, I will not be concerned with morality's overriding obligatory nature. Nor will I assess Korsgaard's answer to the normative question. Rather, I will be concerned with *epistemic* normativity. Just as there are normative claims concerning *how we should act,* it also seems that there are normative claims concerning *how we should believe*. Moreover, not unlike the moral quandary described above, we sometimes confront things that we know to be true, but which are difficult to accept. As it is sometimes said, *the truth hurts.* After all, wouldn't it just be *easier* to believe that human induced climate change isn't real? Or that one's favored political candidate actually *did* win the election? Perhaps the strength of the evidence in certain cases is so clear that we simply *can't help* but believe in accordance with it. Nevertheless, even if this is so, one might still wonder (perhaps, once again, in a Nietzschean spirit), whether we've simply internalized certain psychological tendencies associated with *knowledge*, *truth,* or *justification* (i.e. the subject matter of epistemology) which are downright unhealthy or inhibiting. Thus, it appears that a question akin to Korsgaard's "normative question" can be asked about the claims of epistemology: What accounts for their authority over us? Why should we abide by them? As we

will see, I will focus much of the discussion around the normativity of *evidence*, but I think the discussion will bear in important ways on "epistemic" normativity more generally.

My goal is to vindicate the normative authority of epistemic norms; we have good reason to accept such norms as legitimate standards for the regulation of our doxastic lives. Towards the end of this brief introductory chapter, I will say something about the kind of view that I will develop in order to achieve this goal. However, before getting there, I want to address a worry that one might have regarding the project I'm about to embark on. Just like Korsgaard's dissatisfaction, this worry is a little difficult to make precise. However, I'd like to try to get it on the table. As we will see, I think this worry is misguided. But articulating it will help me frame and motivate the following discussion.

The experience of confronting a difficult moral decision is something that we're all familiar with. We also, in the course of our everyday lives, often face difficult questions about what to believe. Such experiences might lead one to look to philosophy for *guidance*. For instance, it's common to view the philosophical discipline of ethics as a *practical* discipline; it aspires to provide us with *actionable* standards, rules, or principles, i.e. ones that we can *implement* when deciding what to do. Many have thought that this guidance function is somehow importantly tied to the "normative" itself. When it comes to what we should *believe*, there's also a traditional aspiration when it comes to the philosophical discipline of epistemology; viz. that it will provide us with rules, advice, or directions (for simplicity, I'll often just speak of "norms") that we can follow when trying to figure out what to believe. Importantly (and unfortunately), this "traditional" aspiration often gets tangled up with Descartes' foundationalist program in his *Meditations on First Philosophy*. Thus, the traditional aspiration is that the philosophical

discipline known as epistemology will be able to offer an individual guidance *even when* they're entertaining the kinds of skeptical doubts that Descartes was entertaining in the early stages of the *Meditations.*

Setting Descartes to the side for just a second, let's call this general ability of a norm to direct us in the course of our attempts to figure out how to act and think "first-personal guidance". The thought is that, if some norm can be brought to bear on the real, lived circumstances that we find ourselves in, and can actually help us *navigate* these (often difficult) circumstances, then it meets the "first-personal guidance" constraint. Returning now to Descartes, if epistemology is to meet the traditional aspiration described above – if it is to provide us with norms of belief that we can follow *even when* we're entertaining Cartesian doubts – then it must be capable of providing us with norms of belief that meet the "first-personal guidance" constraint *even when we're fully in the grip of such doubts.*

And now here's the (misguided) worry that I want to discuss: Once we've *abandoned* this traditional epistemological aspiration by adopting a broadly "externalist" approach to epistemology, epistemic norms of belief will *fail* to meet the "first-personal guidance" constraint.[2] If epistemic norms of belief which meet this constraint are no longer on the table, then perhaps epistemology doesn't issue in "normative" claims *at all.* For, as I mentioned above, many take "first-personal guidance" to be a distinguishing characteristic of the "normative" itself. If there *are no* epistemic norms of belief that meet this constraint, then maybe epistemology provides us with *no* normative guidance. If this is so, then the project that I'm

[2] Why not, in response, just opt for an "internalist" approach? I'll return to this question momentarily.

about to embark on is perhaps deeply misguided; there is no *epistemic* analogue of Korsgaard's "normative question" since epistemology *just isn't in the business* of providing "normative" guidance at all.

As a first step towards explaining why this worry is misguided, we should ask: Why might someone think that that, once we've adopted a broadly "externalist" approach to epistemology, epistemic norms of belief will fail to meet the "first-personal guidance" constraint? Without going into too much detail, and simplifying matters somewhat, "externalist" approaches to epistemology might be seen as letting go of the traditional aspiration for epistemology described above. Rather than seeking to provide individuals with norms of belief that meet the "first-personal guidance" constraint, epistemologists are in the business of providing *third-personal* theories of, say, knowledge or justified belief. This enterprise need not involve composing a set of rules or norms that agents can implement and use as a guide when deciding what to believe. Moreover, as a theoretical undertaking, it can avail itself of other sound scientific theorizing. Externalists can thus be seen as rejecting Cartesian "first-philosophy" and its associated foundationalist program. Descartes sought an epistemological basis for his beliefs which didn't depend upon *any* of his ordinary commitments. Rather, Descartes relied on nothing more than the bare resources of pure reflection and rational insight. This traditional image of "philosophical" reflection as, all by itself, playing some kind of substantive epistemological role is abandoned by externalists.

We can now perhaps start to see why someone might think that, once we've embraced epistemological externalism, whatever "norms" epistemology has to offer us will fail to meet the

"first-personal guidance" constraint. Externalist epistemologists are no longer attempting to carry out the task that Descartes embarked on; they aren't looking for some distinguishing mark or feature that can be unfailingly apprehended in the course of sustained first-personal reflection. If externalist epistemologists have norms to offer us, then they will likely incorporate *fully objective* criteria, e.g. they might require that our beliefs be *true*, amount to *knowledge*, or be formed by *reliable processes*. However, norms which incorporate the notions of "truth", "knowledge", and "reliability" might be thought to *fail* the "first-personal guidance" constraint. I'll illustrate the point by utilizing the notion of "truth", but one could also utilize either of the other two as well. Consider the following "truth norm" (TN):

TN: Don't believe p unless p is true.

Imagine someone who wants to take up a review of their current stock of beliefs; they want to see whether their beliefs comply with TN. Just knowing the norm and introspecting won't get the person very far when it comes to carrying out this task; for any proposition p that they believe, they must determine whether or not p is true, and TN doesn't tell the person how to do *that.* Because of this, one might think that TN (and other norms that incorporate objective criteria, e.g. knowledge or reliability) fail the "first-personal guidance" constraint; they don't help us *navigate* the difficult terrain that we find ourselves in as believers. As a result of this, one might even deny them the status of "norms" altogether.

This is why one might think that, once we've adopted an externalist approach to epistemology, epistemic norms of belief will fail to satisfy the "first-personal guidance" constraint. One might take the above as a reason to retreat to epistemological "internalism". An internalist might be motivated by the above considerations to formulate norms that incorporate

subjective criteria, e.g. that our beliefs cohere, or somehow align with our current subjective experiences, etc. However, setting aside certain other worries[3], this retreat to internalism simply isn't justified on the basis of the above considerations; externalist norms of belief do *not* fail the "first-personal guidance" constraint. The problem with the above worry is that it assumes that, in order for externalist norms of belief to meet the "first-personal guidance" constraint *in general*, they must meet that constraint *for someone who's fully in the grip of Cartesian doubts.* But this is manifestly false. Consider, for instance, an example I'll return to later[4]. Say that there's a club which is structured by various norms and expectations. One of these norms is that club members should go to club headquarters every Thursday evening at 9 P.M. for a meeting. We can formulate this "meeting norm" (MN) as follows:

MN: If it's Thursday evening, then go to club headquarters for a meeting.

Similar to TN, just knowing MN and introspecting won't be of much help to a person who wants to comply with the norm[5]. But does that really mean the norm can't be *followed*, or that it provides no *guidance*? This is obviously false; a club member who knows MN can follow it quite easily; they can check their phone or a nearby calendar, determine that its Thursday, wait until the evening, and then head to club headquarters. In carrying out these further tasks, they will perhaps be guided by various *other* rules, procedures, and methods, but that doesn't mean

[3] One could raise similar worries regarding norms of belief that incorporate subjective criteria; even for states that are "internal" we can be mistaken about whether or not they obtain. This "anti-Cartesian" view of the mental, while not entirely uncontroversial, has received support (both empirical and philosophical) in recent years: Schwitzgebel (2008), (2011), Williamson (2000, chapter 4). See also: Srinivasan (2015).

[4] Chapter 2.

[5] We can set aside the obvious difference between MN and TN, viz. that compliance with MN requires a physical act (going to club headquarters). Focus, instead, on one's attempt to determine whether or not MN's antecedent is satisfied.

that MN *isn't* a norm. We can put the point this way: Just because, in order to "follow" some norm N, one needs to appeal to various *other* rules, procedures, or methods, that does not mean that N itself provides no guidance or shouldn't be counted as a "norm". The case of MN clearly establishes this.

Return now to our assiduous belief reviewer; the person who's taken up a review of her current beliefs in order to see if they comply with TN. As we noted above, knowing TN and introspecting *alone* will not get this person very far in her task. But hopefully we can now see that, if this person *isn't* in the grip of Cartesian doubts, then she will be fully prepared to appeal to various *other* rules, procedures, and methods in order to follow TN. For instance, if this person considers her belief that her car is currently parked on the corner of Winona Street and Damen Avenue, she can rely on various other methods or procedures to determine the truth of the proposition (she can ask her partner who just walked in, go out and look, etc.). As we've seen, appealing to these further methods and procedures in the course of following TN doesn't mean that TN itself doesn't provide guidance or that it isn't a "norm".

Thus, the adoption of a broadly externalist approach to epistemology doesn't mean that epistemic norms of belief will no longer meet the "first-personal guidance" constraint. This is of course *not* to say that we're immune to error when it comes to following externalist norms of belief – quite the contrary. Despite her best efforts, the assiduous belief reviewer may get things wrong; her appealed to rules, procedures, and methods don't *guarantee* that she'll get things right in every single case. However, this is in lockstep with our broader norm-guided activities. We could say the same thing about the club member who's attempting to follow MN. Perhaps

this person's evidence is misleading; it suggests that it's currently Wednesday when in fact it's Thursday. As a result, he doesn't go to club headquarters and misses the meeting. None of this suggests that externalist norms fail the "first-personal guidance" constraint. It only suggests that agents can sometimes *blamelessly* violate certain norms.

The project that I'm about to embark on is thus not threatened by the above worry. However, the failure of the above worry might open up another concern, and perhaps this one is more pressing. If all that normative epistemology has to offer those of us who aren't in the grip of Cartesian doubts are highly general injunctions like "You should believe the truth!" or "You should use reliable belief-forming methods!", then one might feel pretty let down. Think once more of our assiduous belief reviewer. Perhaps this person was hoping that epistemologists could help her complete her task. If all that's offered to her in response are general injunctions like the ones above, then she'll probably be pretty disappointed. She'll probably say: "Well, I already knew *that*! I was hoping for something more!".

It's here where externalist epistemologists will perhaps split into two camps. On the one hand, you have those who will take the above as reason to, once again, reject the very idea that the philosophical discipline of epistemology is in the "guidance" game. Members of this camp will say that epistemologists are in the business of constructing third-personal theories, and the success of such theories isn't beholden to first-personal adequacy constraints. On the other hand, you have those who will take the philosophical discipline of epistemology to be directly involved in the project of formulating and refining the various "lower level" rules, methods, and procedures, i.e. the ones that our assiduous belief reviewer appealed to in her attempt to comply with TN. As externalists, they will take this to be a philosophical-cum-scientific task; one which

is not constrained by the overly demanding limitations imposed by Descartes.

I will not attempt to adjudicate between these two camps; perhaps their disagreement is illusory, I'm not sure. The point that I want to make is that, by focusing on the debate between these two camps, we have let a certain question slip through the cracks. The fact is that, most people are *not* like our assiduous belief reviewer. Moreover, most people aren't practicing scientists, nor are they philosophers. Imagine, for instance, someone who simply *doesn't accept* knowledge, truth, or reliability as a legitimate standard for the guidance of belief at all; it's better to believe in ways that make one happy, or that solidify communal bonds, or that satisfy the boss, or whatever. What would each of the two camps above have to say to such a person?

Members of the first camp, qua epistemologists, will perhaps have *nothing* to say to this person. They might be interested to see whether this person meets the standards for knowledge or justified belief that are captured in their theories. Since they are externalists, meeting these standards is perhaps not ruled out by accepting one of the "alternative" standards for the guidance of belief mentioned above (e.g the person might, in spite of their *accepted* standards, sill have beliefs that are caused by reliable belief-forming processes). However, their interest in this person, qua epistemologists, will perhaps end there.

What about the second camp? Members of the second camp will perhaps be content if they are successful in their attempts to formulate and refine the various rules, methods, and procedures that can be utilized in the course of carrying out one's investigations and inquiries. Say that they are successful in this task; they've managed to generate an impeccable set of rules, methods, and procedures that can be utilized in the course of various investigations and inquiries.

If the members of the second camp try to offer up the fruits of their labor to the person described above, they will be met with a shrug of the shoulders. After all, this person has an *antecedent aversion or indifference* to truth, knowledge, and justification. Even if the set of rules, methods, and procedures is *impeccable,* this person will respond with little or no interest ("Thanks, but no thanks").

In other words, neither of the two camps mentioned above is directly engaged with the following question: *Why should we accept epistemic norms as legitimate standards for the guidance of belief in the first place?* By "we" I do *not* mean to restrict myself to philosophers and scientists; these are people who have *already accepted* such standards. Philosophers sometimes balk at the very suggestion that someone might accept one of the "alternative" standards for belief mentioned above. As a lover of truth and wisdom, I count myself as someone who rejects these "alternative" standards, but it is simply not my personal experience that, outside of academia, there aren't people who do (this is not even to mention our contemporary "post-truth" political culture.). Also, I recognize that my fictionalized characters (the assiduous belief reviewer and the person who accepts one of the "alternative" standards for belief) are just that: fictions. Matters are complicated and most people probably fall somewhere on a spectrum in between opposing poles occupied by my two characters. However, and once again, it seems that there will be many people who fall on a point on this spectrum closer to the person who accepts one of the "alternative" standards than to the assiduous belief reviewer.

This brings us back to Korsgaard. Asking and answering the kind of question above – a question about *which* norms of belief we should accept and *why* we should accept them – is a *first-personal* question. As a result, it calls for a first-personal *answer*, and this means that it calls

for first-personal *guidance.* As Korsgaard says regarding her normative question,

> the answer we need is really [a] first-personal answer, the one that satisfies *us* when we *ourselves* ask the normative question.[6]

By focusing on the debate between the two camps mentioned above, we've lost sight of this kind of question when it comes to norms of belief. We've either *let go* of philosophy's aspiration to provide guidance, or we've presupposed a certain kind of audience who will *already* be receptive to the kind of guidance that we have to offer. This, I think, is a mistake. If we can provide an answer to the above question, then we will have offered first-personal guidance of a certain kind; guidance vis-à-vis the question of which norms of belief we ought to accept, and why. Moreover, if we don't presuppose a narrow audience, then our guidance will have wide application; it will apply to many who have not *already* accepted epistemic norms.

But perhaps you're thinking: "So what if there are people who haven't already accepted epistemic norms? We don't need to take them into account when doing philosophy; we can get along just fine without giving them any thought". I think this is a bad response, and for many reasons. I'll mention just two. The first reason anticipates an argument that I will be giving later. One thing that I will be arguing for in this dissertation is that we *hold each other* to certain normative expectations qua believers. In particular, I will argue that we hold individuals accountable for complying with epistemic norms *even when* (perhaps *especially* when) these individuals haven't "accepted" these norms. If we're *holding* S to some norm N, then that seems to imply that we think that there's *good reason for S to accept N*. Thus, it seems that we need an account of what these reasons are. What are the reasons that ground the claim that people in

[6] Korsgaard (1996) pp. 16-17.

general should accept epistemic norms? Once again, the existence of such reasons seems to be implied by our social epistemic practice of holding each other accountable for complying with epistemic norms.

Second, and I think even more importantly, the above response suggests a certain view of philosophical inquiry that I think is deeply mistaken. The response suggests that philosophical inquiry is a rather insular affair taken up by trained professionals with a specialized skill set. Perhaps this is the image that is encouraged by the academic *profession* of philosophy, but this is not what *philosophy* should be. Philosophy is a *humanistic* discipline. It concerns itself with (among other things) questions of meaning, significance, and purpose. A first-personal question about which normative standards to accept is a (somewhat stilted) *existential* question; a question about what to commit ourselves to, and why. This is a human question, and as philosophers, we should be engaging with such questions, head on and without apology.

I'll end these preliminary remarks here. I hope the above does enough to motivate and clear space for the project that I'm about to engage in. Before providing brief summaries of the subsequent chapters, I'd like to say something about the kind of view that I will go on to develop. Although I have relied on Korsgaard to help me frame my project, the view that I will end up defending is a far cry from anything that Korsgaard would be willing to endorse (or any other philosopher of a deeply rationalist bent, for that matter). I think the first-personal question of which norms of belief to accept is best answered by attending to our existing social practices. I will not do much by way of attempting to get "underneath" such practices; I won't provide an "external" justification for them, trace their (actual or hypothetical) origins, or formulate some

end, value, or goal that they conduce to. Rather, according to my view, the mere (contingent) *existence* of a certain kind of social practice, and our status as *participants within it*, is what grounds our acceptance of epistemic norms of belief. This is what I intend to argue.

2. Chapter Summaries

The material in this dissertation can be divided into two parts. Chapters 1-3 present arguments against alternative views and the lay the groundwork for the rest of the discussion. Chapters 4-6 develop and defend my positive proposal. Below I provide brief summaries of each of the chapters.

Chapter 1: Against Evidential Minimalism

Evidence is often taken to be "normative" for doxastic agents. What accounts for the normativity of evidence? According to the view that I'll call "evidential minimalism", there is a close connection between strong evidence for the truth of p and a normative reason to believe p: evidence is either *itself* a normative reason for belief, or evidence *gives rise* to such a reason when certain other minimal conditions are met. In this chapter, I argue against evidential minimalism. I argue that there are cases where: (i) an individual S possesses strong evidence E for the truth of p at time t, (ii) all other minimal conditions for the normativity of E are met at t, (iii) S doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at t. I then formulate a plausible linking claim connecting openness to criticism and the existence of normative reasons for belief. I argue that the minimalist can either accept or reject this linking claim. I argue that, either way, the minimalist view falters.

Chapter 2: *Evidence and Epistemic Norms*

Even if the minimalist view is mistaken, there still might be correct norms of belief that

incorporate the concept of evidence. The task going forward will be to understand evidential norms of belief given the failure of the minimalist view. I start this chapter by arguing that the "aim of belief" hypothesis will fail to adequately address the problem of evidential normativity. I then formulate a candidate norm EN with the intention of adequately capturing the normative force of evidence. I go on to relate evidential norms of belief to a broader class of "epistemic" norms. The question, then, is whether or not EN is a correct *epistemic* norm of belief. I argue that there's a prima facie presumption that epistemic norms have a wide scope of "application", i.e. they are binding on a fairly wide range of individuals. I formulate the following "Application-Acceptance Connection" (AAC): If a norm N applies to S, then there's good reason for S to accept N. Thus, one intuitive desideratum for N's being a correct *epistemic* norm is that many people will have good reason to accept N. I go on to distinguish two different questions that we can ask about epistemic norms: the "content" question and the "justificatory" question. I suggest beginning with the justificatory question vis-à-vis epistemic norms in general. I motivate this approach while also noting a significant roadblock confronting it. I go on to sketch how I will overcome the roadblock. I end the chapter by considering, and rejecting, a recent proposal by Jane Friedman which would have us understand epistemic norms as "zetetic" norms, which pertains to the activity of inquiry.

Chapter 3: Against Epistemic Instrumentalism

In this chapter I criticize a number of different ways of pursuing an "instrumentalist" approach to epistemic normativity. Instrumentalists argue that our acceptance of epistemic norms is justified given the instrumental efficacy of such norms. Specifically, instrumentalists argue that we should accept epistemic norms given that doing so represents an effective way to satisfy

our aims, interests, and goals. I consider four different ways of developing the instrumentalist position. I argue that all of these views are ultimately unsuccessful. A recurring issue is that none of the main varieties of epistemic instrumentalist is capable of accommodating the evidential norm EN. Specifically, none of the main varieties of instrumentalism has the resources to secure a wide scope of application for the evidential norm EN. However, at this point in the discussion, EN has yet to be established as a correct epistemic norm. Thus, my full argument against instrumentalism will have to await an argument that doesn't occur until Chapter 6. Nevertheless, insofar as EN captures our commonsense views regarding epistemic reasons for belief, the instrumentalist's inability to accommodate it gives us at least prima facie reason to reject her view.

Chapter 4: Epistemic Accountability

In this chapter I begin to develop my positive account of epistemic normativity. In Chapter 2, I suggested beginning with the justificatory question vis-à-vis epistemic norms in general. However, there's a significant roadblock confronting this approach: There doesn't seem to be a non-arbitrary way of demarcating the class of "epistemic" norms. The way I propose overcoming this challenge is by attending to our actual social practices of interpersonal accountability. Specifically, I propose that we try to discern a *distinctive form of response* that can be taken up by others in reaction to a person's violation of a norm of the relevant kind. This chapter is devoted to establishing that *there is* a form of response like this at play in our actual social practices. I call this response "epistemic accountability". In order to frame the discussion, I utilize the notions of accountability, attributability, and answerability which are often invoked in the literature on practical agency and moral responsibility. I formulate three criteria that a

response R must meet in order to count as an epistemic accountability response. I then argue that there is a form of response at play in our actual practices which meets these criteria.

Chapter 5: Epistemic Accountability as a Legitimate Social Practice

Given the results of Chapter 4, we have a possible way of demarcating the class of "epistemic" norms, as well as a potential answer to the justificatory question vis-à-vis epistemic norms in general: "Epistemic" norms could be thought of those which structure our social practice of epistemic accountability. The "justificatory" question vis-à-vis epistemic norms asks about the reasons which ground our *acceptance* of such norms. Thus, we could potentially answer the justificatory question in the following manner: We should accept epistemic norms given that these norms structure a social practice that we participate in. However, this answer to the justificatory question will only work if our social practice of epistemic accountability is a *legitimate* social practice. In this chapter, I argue that our social practice of epistemic accountability is, indeed, a legitimate social practice. I do this by considering a number of forceful attempts to undermine the legitimacy of the practice. I argue that none of these challenges succeeds. I first consider a challenge concerning instances of epistemic injustice, I then consider a number of challenges related to holding non-culpable believers epistemically accountable.

Chapter 6: Answering the Justificatory and Content Questions

Having established the legitimacy of our practice of epistemic accountability, I go on in this chapter to develop answers to the justificatory and content questions. The answer to the justificatory question reveals our reasons for accepting epistemic norms. The answer to the content question tells us what the norms *are* which structure our social practice of epistemic

accountability. According to my answer to the justificatory question, we have reason to accept epistemic norms given that such norms structure a legitimate social practice that we participate in. According to this answer, it is not the case that the relevant norms bind us in virtue of the fact that there's some independent reason that grounds our acceptance of them. Rather, according to my view, we have reason to accept the norms *in virtue of the fact that they are binding on us*. They are binding on us, once again, given that they structure a legitimate social practice that we participate in. According to my answer to the content question, the evidential norm EN is one norm that structures our social practice of epistemic accountability. Given the answer to the justificatory question, a very wide range of individuals will have good reason to accept EN. This completes my argument against the instrumentalist position. I then go on to argue that our social practice of epistemic accountability is structured by further norms of belief which require not just *true* belief but *knowledge*

Chapter 1: Against Evidential Minimalism[7]

1. Introduction

Strong evidence for the truth of p is often taken to give rise to an epistemic reason to believe p. Evidence is thus commonly taken to be "normative" for doxastic agents; it is the kind of thing that we ought to be sensitive to when forming and maintaining attitudes like belief, disbelief, and suspension of judgment[8]. But what accounts for the normativity of evidence? Is it the case that evidence *itself* is normative for belief? Or is the normativity of evidence accounted for in terms of some *extra*-evidential consideration(s)?

In the following chapter, I will consider, and reject, the view that says evidence *itself* (perhaps when coupled with certain other "minimal" conditions) is normative for belief. I will call this view "evidential minimalism". I will take the minimalist as offering us an answer to the following question: Under what conditions does strong evidence for the truth of p generate an epistemic reason for S to believe p? According to the minimalist's answer, there are either very *few* or even *no* conditions that need to be met, over and above the obtaining of strong evidence for the truth of p, in order for there to be an epistemic reason for S to believe p[9].

---

[7] The following chapter is a lightly revised version of my (Forthcoming). I would like to thank an anonymous reviewer at the journal *Episteme* for helpful feedback and suggestions.

[8] The 'reason' in 'epistemic reason' should thus be understood in the *normative* sense. Normative reasons are considerations that *justify, call for,* or *recommend* attitudes and actions, at least in a pro tanto manner. An "epistemic" reason for belief would be one that supports adopting or sustaining a belief *in a particular way*, e.g. by bearing positively on the belief's standing as knowledge.

[9] Note that one could be an evidential minimalist without also subscribing to views that are labeled as "evidentialism". There are at least two different views that get called "evidentialism". First, there's a view which says that positive epistemological status (e.g. being epistemically justified in believing p) must be tied in some way to evidence. Thus, on this view, one can't be epistemically justified in believing p unless one has evidence for p. Second, there's a view which says that there *aren't* "pragmatic" reasons for belief; only considerations bearing on the truth of propositions (i.e. evidential considerations) constitute reasons for belief. These views are distinct (e.g. one could subscribe to the former but not the latter by holding that there are pragmatic

For instance, according to one way of being a minimalist, the obtaining of strong evidence is *itself* enough to generate an epistemic reason for S to believe p. I'll call this version of minimalism "objectivism" since it doesn't require that S *possess* or otherwise *attend to* the evidence in order for it to give rise to an epistemic reason for S to believe p. On another way of being a minimalist, there are certain minimal conditions that need to be met – over and above the obtaining of strong evidence for the truth of p – in order for there to be an epistemic reason for S to believe p. For instance, perhaps S has to possess the evidence, or possess the evidence and consider the question of whether p. Why think of these latter conditions as "minimal"? The idea is that "possessing" evidence and "considering whether p" are conditions that can be met rather easily; once strong evidence for p is already on the table, a normative reason to believe p isn't too far off, so to speak. Any normal epistemic agent possesses all kinds of excellent evidence at any given time, and questions of the form "whether p" can be thrust upon our conscious awareness from without (e.g. by an annoying friend or co-worker).

> Consider, for instance, the view put forward in Hofmann (2021) according to which:
>
> neither a bit of evidence nor the fact that it is evidence for a certain proposition *is* a normative fact, but it is still the case that evidence *provides* normative reason for belief. In this latter sense, then, evidence is normative.[10]

Hofmann argues that, in the prospective case (where a subject S possesses strong evidence for p but doesn't yet believe p), certain conditions have to be met in order for there to be a normative claim upon S to believe p. For instance, it has to be the case that "the question whether p is *activated* for S and S has all the abilities needed for following [the evidence]"[11]. Similarly,

---

reasons for belief, they're just not *epistemic* reasons for belief). However, commitment to evidential minimalism doesn't imply commitment to either of these views.

[10] Hofmann (2021) p. 667.

[11] Ibid. p. 678 (emphasis added).

Kiesewetter (2017) argues that "If A *has* sufficient evidence for p, and A *attends* to p, then A is rationally required to believe p."[12]. I will count these further conditions – possessing the evidence, "activating" the question of whether p, paying attention to p, etc. – as "minimal". The rough idea is that *evidence* is still shouldering most of the normative weight. Even though there are certain conditions that have to be met, over and above the obtaining of the evidence, in order for there to be a normative reason for S to believe p, these conditions don't add *too* much to the picture. In particular, they don't seem to add anything which would threaten to usurp the central role that evidence is playing in explaining why S has a normative reason to believe p.[13]

The minimalist view thus comes in different shapes and sizes. While the minimalist/non-minimalist distinction will perhaps not always be clear-cut, we can provide a preliminary disjunctive characterization of the minimalist position[14]. Evidential minimalists subscribe to either one of the following:

Strong evidence for the truth of p is *itself* an epistemic reason for S to believe p.

---

[12] Kiesewetter (2017) p. 185 (emphases added).

[13] Commitment to minimalism is sometimes expressed in opposition to certain *instrumentalist* or *teleological* approaches to epistemic normativity. See, for instance: Kelly (2003), (2007), and Berker (2013). According to Kelly (2007), "possession of evidence is itself something which has normative import, and…to possess strong evidence that some proposition is true is *ipso facto* to have reason to think that that proposition is true." (p. 473).

[14] A note regarding terminology: the view that I'm calling "minimalism" is also sometimes called "normativism" or "intrinsicalism" in the literature. For instance, Kiesewetter (2021) labels as "normativism" the view which says that "epistemic reasons are normative reasons for belief." (p. 2). Similarly, Schmidt (2021) claims that "normativism" is the view which says that "purely evidential considerations provide us with reason for belief." (p. 3). Cowie (2014) labels as "instinsicalism" the view which says that "there is reason to believe in accordance with one's evidence in virtue of a *brutely epistemic* normative truth relating belief to evidence, or to some other epistemic property such as truth or epistemic rationality." (p. 4004). The rough idea, once again, is that *evidence itself* is normative for belief. I will stick with the "minimalist" designation throughout. The minimalist view about epistemic reasons is also sometimes associated with "primitivism" about normative reasons generally (see, for instance: Scanlon (1998), Parfit (2011)).

Or

> Strong evidence for the truth of p *gives rise to* an epistemic reason for S to believe p when certain other minimal conditions are met (e.g. S possesses the evidence and considers whether p).

A *non*-minimalist about evidential normativity, by contrast, would hold that something more substantial needs to be on the table, over and above the obtaining of strong evidence for the truth of p and certain other minimal conditions, in order for there to be an epistemic reason for S to believe p. The recent literature on epistemic normativity provides a number of non-minimalist possibilities. For instance, Steglich-Petersen (2018) argues that evidence for p constitutes a normative reason for S to believe p only when S has some *further* (e.g. moral or prudential) reason to form true beliefs about p. Thus, the evidence *by itself* (even when coupled with certain minimal conditions) isn't enough to underwrite a normative reason for belief[15]. Such a view is also implicit in Papineau (2013) where it is argued that epistemic evaluations carry "no prescriptive force on [their] own, independent of some further value attaching to the aim of truth."[16] Similarly, Maguire and Woods (2020) argue that evidential considerations are not "genuinely" or "authoritatively" normative. According to Maguire and Woods, "all and only the *practical* reasons are the authoritative reasons."[17] Non-minimalism is thus well represented in the contemporary literature.[18] These views are unified as "non-minimalist" insofar as they introduce conditions which *do* threaten to usurp the central role that minimalists seek to reserve for

[15] Nelson (2010) makes a similar claim. According to Nelson, there are no *positive* epistemic duties (i.e. duties *to believe* specific propositions), only *negative* epistemic duties (i.e. duties to *refrain* from believing certain propositions). Whenever there's a positive duty for S to believe p, some *extra*-epistemic (e.g. practical) consideration must be involved.
[16] Papineau (2013) p. 69.
[17] Maguire and Woods (2020) p. 229.
[18] Other proponents include: Cowie (2014), Mantel (2019), McCormick (2020), and Rinard (2015).

evidence when explaining why S has a normative reason to believe p.[19]

Motivation for non-minimalism is sometimes provided by considering the "justificational fecundity" of evidence; the fact that evidence potentially justifies an infinite number of beliefs[20]. Consider, for instance, Whiting (2013):

> Suppose that I have evidence that the cakes are burning. Whatever provides this evidence provides evidence for an infinite number of other beliefs, such as that there are cakes, that the cakes are burning or that Tolstoy wrote *Great Expectations*, that if the cakes are burning then the cakes are burning, that there is more carbon dioxide in the atmosphere than there was before I made the cakes, that there is no dancing monkey singing the national anthem in the space the cakes occupy, and so on without end.[21]

One could take such observations to suggest that, even if one has *decisive* evidence E for the truth of p, one still does not have a normative reason to believe p. It is at precisely this point, however, where the minimalist view stands to be bolstered by appealing to minimal conditions. For instance, it's less clear that Whiting would lack a normative reason to believe some of the claims in the above passage if, in addition to possessing the relevant evidence, certain questions were to be brought before his mind. One could argue that this so even in those cases where Whiting lacks practical reason to be interested in the relevant claims.

Plausible as this might sound, I will argue that it is mistaken: the minimalist view is false. I will not argue against minimalism by defending the comparative theoretical strengths of some version of non-minimalism. Rather, I will argue directly against the minimalist view itself. If my argument is successful, then *some* form of non-minimalism will be the way to go when it comes

---

[19] For instance, if there has to be some *practical* interest at stake, over and above the obtaining of strong evidence for the truth of p and certain other minimal conditions, in order for there to be a normative reason to believe p, then it doesn't look like *evidence* is doing the heavy lifting, normatively speaking.

[20] I borrow the "justificational fecundity" label from Nelson (2010).

[21] Whiting (2013) p. 130.

to understanding the normativity of evidence. In the present chapter, I remain neutral on which form of non-minimalism should be preferred (I will explore this question in subsequent chapters). The bulk of the chapter will be devoted to my argument against minimalism. However, before presenting the argument (sec. 3), I would like to say a few things regarding evidence and its possession (sec. 2).

2. Evidence and its Possession

There are a number of debates concerning evidence in contemporary epistemology[22]. Since it is beyond the scope of the present chapter to argue for one particular approach, I will simply flag how I'll be understanding the notion throughout. Note that while the following assumptions about evidence are contentious, the success of my argument against minimalism doesn't essentially depend upon them. In other words, similar arguments could be offered against minimalism that operate with alternative conceptions of evidence. Nevertheless, it will be helpful to introduce the following simplifying assumptions.

First, I will assume that evidence is *propositional*. This rules out as evidence experiential states that are not propositionally structured, as well as ordinary objects and artefacts, e.g. fossils or murder weapons[23]. What this rules *in* as evidence are propositional claims like "that the victim's blood is on the knife". Second, I will assume that evidence is *factive*. In other words, in order for E to be evidence for the truth of p, E must be *true*. This rules out *false* propositions as evidence[24]. Third, I will assume that our evidence can include true propositions *about the*

[22] For overviews, see: Fratantonio (forthcoming), Kelly (2008), (2014).
[23] Why rule out ordinary objects and artefacts as evidence? One reason is that evidence is often taken to stand in certain relations that obtain between propositions, e.g. relations of probabilistic support.
[24] Williamson (2000) argues for a stronger claim. According to Williamson, all and only *known* propositions are evidence. Leite (2013) argues for the weaker claim that evidence cannot be false. Similarly, Littlejohn (2013) argues that any propositional evidence must be true.

external world. The propositional nature of evidence combined with factivity doesn't automatically guarantee this third assumption. For instance, one could hold that evidence consists of all and only true propositions *about how things appear to us.* I will assume a more expansive view of evidence according to which true propositions about the external world can be a part of one's evidence.[25]

I will also focus on evidence E for the truth of p *of a certain strength.* In order for E to be evidence for the truth of p, E must somehow *indicate* or *make more likely* the truth of p. It's possible, however, for some proposition E to only raise the likelihood of p ever so slightly. To simplify matters, I will focus on instances of evidential support that are very strong. Specifically, I will focus on instances where the support relation between E and p is such that one could come to *know* p on the basis of E.[26] I will not assume that this support relation must amount to *entailment*; I take it that I can know things (say, on the basis of testimony or memory) without my evidence entailing the truth of the relevant propositions. However, entailing evidence would certainly *suffice* to qualify as "strong" evidence. Thus, for the purposes of the following discussion, we can say the following:

> In order for E to be "strong" evidence for the truth of p:
>
> (i) E must meet the constraints on evidence introduced above, i.e. E must be a true proposition (possibly about the external world); and

[25] The above conception of evidence is thus "externalist" in many respects. For instance, it denies that evidence must supervene on a subject's non-factive mental states. It also allows for the possibility that systematically deceived individuals (e.g. brains-in-vats) and their non-deceived, internally indistinguishable, counterparts can *vary* in terms of the evidence that they possess. For some potential worries for this sort of externalist view see: Silins (2005).

[26] I do not mean to suggest that the support relation between E and p must be this strong in order for E to be *evidence* for p. Rather, I am simply focusing the discussion on very strong *instances* of evidential support.

(ii) E must make the truth of p sufficiently likely; it must be possible for one to come to *know* p on the basis of E.

Focusing matters in this way simplifies the discussion in certain respects. For instance, if it can be shown that the minimalist view falters *even when* we restrict our attention to evidence that is "strong", then it doesn't seem that the view holds much promise.

Let's turn now to evidence "possession". Given the broadly externalist view of evidence presented above, it's possible for evidence to obtain without a person's being in possession of it. Also, recall that, according to one way of being a minimalist, strong evidence E for the truth of p isn't *itself* an epistemic reason for S to believe p. Rather, certain other "minimal" conditions must also be met, e.g. S must *possess* the evidence and consider whether p. What is it to "possess" evidence? Let's focus on the following example discussed in Kelly (2007), (2014): Koplik spots are evidence of measles. This is a true claim, and it reflects a discovery of medical science. Say that patient X presents with Koplik spots. Let's also say that the presence of Koplik spots makes the claim that the patient has measles sufficiently likely; one could come to know this claim on the basis of the evidence. We thus have a case of "strong" evidence, with the following evidence E and supported proposition p:

E: Patient X has Koplik spots.

p: Patient X has measles.

When it comes to the question of evidence "possession", we can distinguish between three different types of cases. In the first type of case, a subject S has no contentful mental state (whether occurrent or non-occurrent) which represents E as true. This is a clear case in which a subject doesn't "possess" the evidence. For instance, perhaps the subject has no such mental state because they haven't been made aware of E; they haven't seen the patient, nobody has told them about the patient, they haven't read about the patient, etc. Next there's the case in which a

subject S clearly *does* possess the evidence. Imagine a doctor who is attending to patient X. Say that the doctor is aware of the connection between Koplik spots and measles, is proficient at identifying the spots, and wants to know about the patient's medical condition. After attending to the patient and noticing the spots, the doctor clearly "possesses" the evidence.

Finally, there's an intermediate case. Imagine that there's a different doctor attending to the same patient. However, this doctor is *ignorant* of the connection between Koplik spots and measles. Even though this may be the case, the doctor can still attend to the patient's spots. It's not clear whether, in such a case, the doctor "possesses" evidence that the patient has measles. Here we can simply say that there's a very *weak* sense in which this doctor "possesses" the relevant evidence. The second type of case mentioned above can thus be considered evidence possession in a more robust sense; "robust" possession for short.

We can now give a rough gloss on evidence "possession". The kind of possession that is relevant for the minimalist is evidence possession in the *robust* sense. We can thus say the following:

> A subject S possesses strong evidence E for the truth of p in the "robust" sense at time t iff (i) S has, at time t, some contentful mental state (whether occurrent or non-occurrent) which represents E as true, and (ii) S is aware of E's evidential import vis-à-vis p at t.[27]

A few clarificatory remarks: In the example of robust possession above, the doctor had an *occurrent* mental state at time t, and was also, at time t, explicitly and consciously appreciating E's evidential import vis-à-vis p. While this will suffice for "robust" possession, it's not necessary. For instance, the mental state which represents E as true doesn't have to be *occurrent.* An individual could, for instance, simply *know* that E obtains; something which could be

[27] Similar to my assumptions about evidence, this construal of evidence possession should be taken as a rough and ready conception that will help facilitate the following discussion.

committed to memory and then "called up" if the individual were to consider the matter. Also, "awareness" of E's evidential import vis-à-vis p is an *ability* or *capacity* had by S; something which S can have, at some time t, even if she's not, at that time, explicitly or consciously appreciating E's evidential import vis-à-vis p. Thus, an individual S can, at time t, possess (in the "robust" sense) strong evidence E for the truth of p without attending to E at t *and* without believing p at t. From here on out when I speak of evidence "possession" I will have in mind possession in the *robust* sense.

3. An Argument Against Minimalism

My argument against minimalism will occur in two stages. In the first step (sec. 3.1) I will establish the following claim:

> There are cases where: (i) a subject S possess strong evidence E for the truth of p at time t, (ii) all other minimal conditions for the normativity of E are met at t, (iii) S doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at time t.

The second step of the argument (sec. 3.2) will involve the following linking claim connecting openness to criticism with epistemic reasons for belief:

> **If** [there are cases where: (i) a subject S possess strong evidence E for the truth of p at time t, (ii) all other minimal conditions for the normativity of E are met at t, (iii) S doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at time t] **then**, in such cases, it's not the case that E is (or gives rise to) an epistemic reason for S to believe p at t.

I think there is some plausibility to this linking claim, and I will say more about that below. However, my argument won't depend upon the truth of this claim. Once the claim in step 1 is established, the minimalist is faced with the following choice: She can either accept or reject the linking claim. I will argue that, either way, the minimalist view falters.

Before proceeding to carry out the first step of the argument, I'd like to say something about condition (iv) "S isn't open to any form of criticism on account of (i)-(iii) at time t".

Specifically, I'd like to clarify what I mean by "criticism". When it comes to "criticism" I mean to be somewhat inclusive. Thus, I do not restrict myself to *blame*. Some authors hold that there is a distinctly *epistemic* form of blame[28]. If such a thing exists, then I would include it as a form of criticism. However, I don't mean to *restrict* myself to blame. Criticism might involve aretaic assessment of the person which falls short of blame. Also, there are ways of holding a person accountable (e.g. via sanctions of various kinds) which do not imply that the person manifests a character defect. What I *do* mean to rule out is something that we might call "mere appraisal"; simply assessing someone's thought or conduct vis-à-vis a standard where this has no real weight or significance in our actual practices. The thought is that, in order to count as "criticism", the form of response at issue must somehow go beyond mere assessment vis-à-vis a standard. For instance, blaming someone is sometimes thought to involve certain reactive emotions or modification of various attitudes and expectations vis-à-vis the person[29], sanctioning someone might involve imposing certain punitive measures, etc. The question of whether there is a *distinctly epistemic* form of response that counts as "criticism" in this sense is an interesting one, and I will return to it below. For now, I'll simply note that "criticism" must go beyond mere assessment in some way.

3.1 An Argument Against Minimalism: Step 1

The first step of my argument against minimalism will involve establishing that there are

[28] Boult (2021a), (2021b), (2021c), (2020), Brown (2018), Schmidt (2021).

[29] A number of philosophers attempt to account for the characteristic "sting" or "force" of moral blame in terms of certain emotional responses that typify ordinary interpersonal relationships. This is a broadly Strawsonian approach to moral blame (Strawson (1962/2003)). Proponents include Wallace (2011), and Wolf (2011). Scanlon (2008) proposes an alternative account of moral blame which downplays the importance of the reactive sentiments while also preserving the centrality of interpersonal relationships. Boult (2020) develops a broadly Scanlonian account of *epistemic* blame. I'll return to Boult's account below.

cases where the above conditions (i)-(iv) are met. Keeping in mind the account of evidence "possession" presented earlier, consider the following case:

> BORED AT HOME: Katlyn is at home with nothing to do. She recently left her job in order to take up a better position at a new company. There's a one month interval in between her leaving the old job and beginning training for the new job. Katlyn doesn't have to move for her new job, and she's tied up all manner of loose ends, taken care of various errands and housekeeping, etc. She finds herself at home one Sunday afternoon and is simply bored with nothing to do.

As a typical epistemic agent, Katlyn finds herself in this situation with a large stock of evidence upon which she could draw in order to arrive at new knowledge. For instance, on this particular Sunday, Katlyn could sit down and start drawing out consequences from things that she already knows, e.g. she could try to figure out how many hours she's been alive, or how many haircuts she's received in her life so far. Also, while her house will certainly not be as lively as, say, a busy city street, there will still be quite a bit of perceptual information that she could attend to in order to acquire new knowledge, e.g. she could figure out exactly how many tiles there are on her bathroom floor. However, Katlyn of course does not do this. Nor do we expect her to. We wouldn't think any less of Katlyn for not doing this. In fact, we'd probably be somewhat taken aback if she *did* start doing this.

Imagine one of the questions I mentioned above: How many hours has Katlyn been alive? Say that Katlyn knows her own age (32), and she's quite talented at quick mental math. She could easily figure out that she's been alive for more than 200,000 hours, were she to sit down and think about it for a minute. Thus, at any arbitrary time t when Katlyn is at home, she will possess, in the "robust" sense, strong evidence E for truth of the following claim: that she has been alive for more than 200,000 hours. However, Katlyn is of course in no way *criticizable* for not believing this claim. At any arbitrary time t when Katlyn is bored at home not thinking about the question of how many hours she's been alive, she manifests no fault and opens herself up to

no form of blame, sanction, or criticism for *not* believing this claim. Examples like this are easy to generate. Here's another thing that Katlyn knows: that her name is Katlyn. Katlyn also remembers the disjunction introduction rule from her symbolic logic class in college. Thus, at any arbitrary time t, Katlyn will possess, in the "robust" sense, strong evidence for the following disjunctive claim (DC):

> DC: Either her name is Katlyn or aliens exist.

As well as:

> DC*: Either her name is Katlyn or aliens exist or the moon is made of cheese.

And:

> DC**: Either her name is Katlyn or aliens exist or the moon is made of cheese or Santa Claus is real.

And so on. These propositions are, quite clearly, a bunch of disjunctive junk. Once again, Katlyn manifests no fault and opens herself up to no form of blame, sanction, or criticism, for failing to believe them. Consider one further case:

> LAST DONUT: Someone at work ate the last donut. The box of donuts was in the lunchroom. David knows an assortment of information which is such that, were he to reason through it properly, he could determine who ate the last donut. For instance, he knows that Sarah always eats lunch in the lunchroom, but only between 12 and 12:45, that the last donut was present until approximately 1:15, that Sue was off that particular day, that Jamie doesn't like donuts, etc. Say that Carl ate the last donut. Moreover, Carl wasn't violating office protocol in doing so. Actually, David himself brought the donuts and placed a note on the box that said "please eat me".

Say that, at time t, David knows everything he needs to know in order to figure out who ate the last donut and that figuring this out would take some concentrated effort but wouldn't be exceedingly difficult; David is fully capable of figuring it out. However, David takes no interest in this; he *wanted* the donuts to be eaten. Similar to Katlyn, David is in no way criticizable, at t, for not believing that Carl ate the last donut; he manifests no fault and opens himself up to no

form of blame, sanction, or criticism for not believing this claim.

The cases of BORED AT HOME and LAST DONUT perhaps suffice to establish the following claim:

> There are cases where: (i) a subject S possesses strong evidence E for the truth of p at time t, (ii) S doesn't believe p at t, yet (iii) S isn't open to any form of criticism on account of (i)-(ii) at time t.

However, this was not the claim we were to establish. That claim was as follows:

> There are cases where: (i) a subject S possess strong evidence E for the truth of p at time t, (ii) *all other minimal conditions for the normativity of E are met at t,* (iii) S doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at time t.

A minimalist could happily grant that BORED AT HOME and LAST DONUT suffice to establish the first claim but not the second. This minimalist could argue that there are further minimal conditions for the normativity of evidence that aren't yet met in the relevant cases. Let us, then, introduce further minimal conditions for the normativity of evidence. As I mentioned earlier, these will include things like being prompted with questions of the form "whether p" and attending to the relevant evidence. It is important to note that there will be certain *non*-minimal conditions that we will want to exclude. Imagine, for instance, BORED AT HOME*. In BORED AT HOME* everything is the same except that Katlyn becomes genuinely interested in the question of how many hours she's been alive. The question is before her mind and she wants to figure out an answer to it. She then "activates" her mathematical abilities, engages in a process of reasoning, and arrives at an answer to the question. Katlyn is thus engaged in a certain goal-directed or purposive activity; she's engaged in active inquiry when it comes to a certain question or subject matter. We can also imagine David doing so in similarly adjusted LAST DONUT*. Engagement in such an activity would be a *non*-minimal condition; it is precisely the kind of thing that threatens to usurp the central role that minimalists wish to reserve for evidence

in explaining why S has an epistemic reason to believe p. Consider, for instance, the following example adapted from a case discussed in Kelly (2003):

> MOVIE SPOILER: Liz often sees newly released movies only after they've been in theaters for some time. Prior to her seeing some particular movie, she has no interest in forming true beliefs about its ending. In fact, she actively does *not* want to take on such beliefs. Liz is chatting with a few friends before class. One of them quickly changes topic and begins talking about a movie that Liz hasn't yet seen. Before Liz has a chance to warn her not to spoil the ending the friend blurts out the ending and Liz comes to form a true belief about its ending.

In his (2003), Kelly argues that our epistemic reasons for belief do *not* depend upon our adoption of certain aims or goals, including the goals that we take on when engaged in active inquiry vis-à-vis some topic or subject matter. Regarding the movie example, Kelly notes that, when it comes to Liz's epistemic reasons for belief, there would be no difference between MOVIE SPOILER and MOVIE INFORMER; in the latter case everything is the same except now Liz *wants* to know about the ending of the movie and she *asks* her friend to divulge this information. According to Kelly, the goals and desires had by Liz in MOVIE INFORMER are extraneous and irrelevant when it comes to explaining what her epistemic reason for belief are.

Thus, we shouldn't count something like "S is engaged in active inquiry vis-à-vis whether p" as a minimal condition for the normativity of evidence; this would threaten to obscure the central role that evidence is playing in the minimalist's explanation as to why S has an epistemic reason to believe p. Somewhat ironically, then, the normativity of evidence seems to be most clearly on display in cases in which a subject is *not* engaged in conscious deliberation or active inquiry vis-à-vis some subject matter or question. But what, then, *would* count as a minimal condition for the normativity of evidence? As I've already mentioned, simply prompting S with a question of the form "whether p" will perhaps be enough, as will getting S to attend to the relevant evidence.

Let's return again to BORED AT HOME and LAST DONUT. Say that Katlyn has a friend over and her friend turns to her and says "Hey, you're 32 and good with math, how many hours have you been alive?". Imagine that, in response, Katlyn simply bats the question away, refusing to "activate" her mathematical abilities because the question is silly and pointless. She thus doesn't go through the requisite process of reasoning in order to arrive at an answer (one that she easily *could* go through). Once again, it doesn't seem that Katlyn is open to criticism on account of this. She manifests no fault and opens herself up to no form of blame, sanction, or criticism for responding in the way that she did. Similarly, in LAST DONUT, a coworker might approach David and ask "Who ate the last donut?". Even though David is in possession of everything he needs in order to answer this question, and *could* answer the question after giving it a little thought, he is in no way criticizable for simply answering "I don't know".

Thus, even if certain minimal conditions for the normativity of evidence are met, a subject can still fail to be appropriately subject to criticism on account of not believing the target proposition. I think this lends strong support for our claim that there are cases where conditions (i)-(iv) are met. However, one might remain unconvinced. One could argue that, in the modified version of BORED AT HOME (where Katlyn's friend asks her how many hours she's been alive), it is not the case that *all* of the minimal conditions for the normativity of evidence are met[30]. For instance, let's say that Katlyn knows the following conjunctive proposition E:

E: She is 32 years old & there are 24 hours in a day & there are 365 days in a year.

Say that, at time t, Katlyn is attending to E and she is asked "How many hours have you been alive?". Perhaps, at time t, there are certain minimal conditions for the normativity of E that remain unmet. For instance, perhaps Katlyn has to perform the relevant calculations and thereby

[30] I'd like to thank an anonymous reviewer from the journal *Episteme* for raising this objection.

come to consciously and vividly "see" the evidential connection that obtains between E and the target proposition. Consider the following analogy: Say that someone has to look around the corner in order to acquire some evidence E bearing on the question of whether p. Prior to looking around the corner, perhaps the evidence doesn't give rise to a normative claim regarding what the person should believe. One could argue that Katlyn's case is similar. In Katlyn's case, the relevant "act" is, of course, mental rather than physical. However, perhaps a similar point holds: prior to engaging in this act, the evidence E generates no normative claim concerning what Katlyn should believe.

The problem with this proposal is that there are cases where conditions (i)-(iv) are met *and* S is "consciously and vividly" appreciating the evidential connection that obtains between E and the target proposition at t.[31] Return again to the disjunctive claims we saw earlier in the case of Katlyn in BORED AT HOME. Say that Katlyn consciously acknowledges that she knows her name is Katlyn at time t. Say that she also consciously and vividly and acknowledges that the following disjunctive claim (DC) follows from this known fact:

DC: Either her name is Katlyn or aliens exist.

If the question of "Whether or not DC?" is somehow brought to her attention, would Katlyn *then* be appropriately subjected to criticism were she to not believe DC? It's hard to imagine Katlyn being more "in touch" with the relevant evidence. Perhaps we might expect Katlyn to *assent* to DC in such a circumstance. However, if Katlyn could somehow avoid *believing* DC even after

[31] To be clear, I'm not conceding that "consciously and vividly" appreciating E's evidential import vis-à-vis p is a minimal condition for the normativity of evidence; I'm inclined to think that it isn't. Rather, I'm claiming that, even if it *were* a minimal condition, that still wouldn't save the minimalist since there are many cases where that condition is met – along with the other minimal conditions already mentioned – yet a subject can still fail to be appropriately subjected to criticism on account of not believing the target proposition.

assenting to it (maybe immediately after assenting to DC, the claim simply evaporates from her conscious awareness), would she thereby be open to some form of blame, sanction, or criticism? Paakkunainen (2018) suggests that she would be:

> there is a clear sense in which, if S has excellent, undefeated epistemic reasons to believe that p, then S ought to believe that p: he epistemically ought to believe that p. At least, this is so if S also considers the reasons, q, and the question whether p…If S considers the question whether there is an even number of dust specks on his desk, and considers the excellent evidence for answering "yes," then he epistemically ought to believe that there is an even number of dust specks on his desk. Likewise, we (well, I) *would* fault him for failing to form that belief if he considered the question and his evidence was clear. He's not *merely* lacking a psychological compulsion if he fails to form the belief in this instance; he's being a less-than-excellent epistemic agent.[32]

If we apply Paakkunainen's considerations to the case of Katlyn and DC do we get the same result? In other words, would Katlyn be a "less-than-excellent epistemic agent" and an appropriate target of criticism is she were to fail to believe DC? My own reaction is that she would not. If Katlyn were to fail to believe DC she would not be appropriately subjected to any form of criticism. Nor would she be if she failed to believe DC* in similar circumstances:

> DC*: Either her name is Katlyn or aliens exist or the moon is made of cheese.

Or DC**:

> DC**: Either her name is Katlyn or aliens exist or the moon is made of cheese or Santa Claus is real.

Let us, however, explore this a little further. In order to do so, I suggest we consider a few recent attempts to pinpoint a *distinctly epistemic* form of assessment that counts as "criticism". If there is such a form of assessment – and I'm inclined to think that there is[33] – then we can ask about

[32] Paakkunainen (2018) p. 135-36.

[33] I thus disagree with Schmidt (2021) who suggests that the existence of a distinctly epistemic form of criticism supports evidential minimalism. According to Schmidt, "the reactive attitudes within our epistemic practice [of interpersonal criticism] reveal the normative significance of

the conditions under which it is appropriately taken up towards a person. If Katlyn *is* appropriately subjected to some form of criticism on account of her failure to believe the disjunctive claims (as Paakkunainen's remarks suggest), then it seems that a *distinctly epistemic* form of criticism would perhaps be our best candidate. The recent literature on this issue is often framed in terms of the concept of *blame*[34]. I will (somewhat hesitantly[35]) follow suit.

Boult (2021a) identifies four recent proposals when it comes to the nature of epistemic blame. First, there is the *emotion-based view*[36]. According to this view, "epistemic blame is the manifestation of reactive attitudes such as indignation and resentment, directed towards a target as a result of the judgment that the target has (culpably) violated some epistemic norm."[37] If epistemic blame is understood in this way, then Katlyn is clearly not appropriately subjected to it; she is not the appropriate target of *resentment* or *indignation* on account of her failing to believe pointless disjunctive claims. Responding to her in this way would be wildly overblown and uncalled for.

Second, there's the *desire-based* view[38]. According to the desire-based view, epistemic blame consists in a characteristic set of dispositions (e.g. to reproach, feel upset, and to verbally request reasons) associated with a certain belief-desire pair. The relevant belief-desire pair concerns "believing badly"[39], i.e. the belief-desire pair concerns "those whose *beliefs* violate

---

purely evidential considerations." (p. 3, emphasis added). I agree that there is a distinctly epistemic form of criticism, but I don't think that this supports evidential minimalism.

[34] For a helpful overview of some of this recent literature see: Boult (2021a)

[35] I prefer the wider notion of "accountability" over "blame"; one can be *held accountable* for thinking or doing something without also being *blameworthy* for thinking or doing that thing. I'll return to this issue in Chapter 4.

[36] Boult cites as adherents: McHugh (2012), Nottelmann (2007), and Rettler (2018).

[37] Boult (2021a) p. 5.

[38] For a defense of this view see: Brown (2020). Brown draws on the work of Sher (2006).

[39] Brown (2020) p. 399.

epistemic norms without excuse."[40] On this account, epistemic blame is what occurs when a person's desire that someone else *not* culpably violate an epistemic norm (e.g. by believing badly) is frustrated. If we understand epistemic blame in this way, then Katlyn is not appropriately subject to it. First, Katlyn doesn't *believe* badly; her (alleged) epistemic infraction is that she *failed to form* certain beliefs. Second, people do not harbor the kind of desire that would be needed in order to make epistemic blame appropriate in the case of Katlyn. Our desires aren't frustrated when individuals fail to believe pointless and silly disjunctive claims.[41]

Third, there's Boult's own *relationship-based* view of epistemic blame[42]. According to Boult's view, epistemic blame consists in a kind of *relationship-modification*[43]. Boult argues that, as members of epistemic communities, we stand in "epistemic relationships" with one another. When a person A epistemically blames B, A judges that B "has done something that falls short of the "normative ideal" of this relationship"[44] and A modifies her attitudes and expectations vis-à-vis B accordingly. For instance, imagine that A finds out that B is dogmatic and biased when it comes to a certain topic or issue. In response, A might cease to trust B's words when it comes to this topic. If we understand epistemic blame in this way, then Katlyn is, once again, not appropriately subjected to this form of response. In not believing pointless

---

[40] Ibid. (emphasis added).

[41] An anonymous reviewer from the journal *Episteme* has suggested to me that people generally desire that the members in their epistemic community are rational, and that Katlyn frustrates this general desire by not believing the pointless disjuncts. Even if people generally hold such a desire, I find it highly implausible that Katlyn's behavior would frustrate it; I, for one, wouldn't care at all if someone in Katlyn's situation failed to believe pointless disjunctive claims. McCormick (2020) makes a similar point when discussing the desire-based view: "I think we have this desire [that people not believe badly] when the ignoring and flouting [of epistemic reasons] *matters* or when we view it as *mattering*." (p. 43).

[42] Boult (2020), (2021c).

[43] Boult draws on Scanlon's account of moral blame in developing his view: Scanlon (2008), (2013).

[44] Boult (2021a) p. 6.

disjunctive claims, Katlyn is not doing something which compromises her epistemic relationship with other people. Matters would perhaps be different if Katlyn had a *general tendency* to not believe the claims that are well supported by her evidence (say, out of dogmatism or wishful thinking). However, this is not what's going on in her case; Katlyn is simply avoiding taking on unnecessary and pointless beliefs. Note that her failure to form the relevant beliefs in this case might still express an underlying tendency or disposition, viz. *to avoid believing pointless and trivial claims.* However, as long as this tendency is only restricted to those truths that are pointless and trivial, then Katlyn is not open to any form of criticism on account of her being disposed in this way.[45,46]

Finally, there's what Boult calls the *agency-cultivation* view of epistemic blame.[47] This is a broadly "forward-looking" account of epistemic blame which says that blame functions as a "vector" for agency cultivation. According to this view, blame-responses function so as to discourage certain forms of behavior. By internalizing the expectations and demands implicit in the blame-responses of members of our communities, we become responsive to certain salient

[45] Schmidt (2021) builds on Boult's account and says that epistemic blame occurs when we modify epistemic trust in response to a person's *vice* (e.g. gullibility, dogmatism, or wishful thinking). However, as I've indicated, Katlyn's failure to take on pointless disjunctive beliefs is not a manifestation of a vice. In fact, Katlyn's tendency might even constitute an intellectual *virtue*; she expends her cognitive resources wisely by avoiding cluttering her mind with useless and pointless junk. Cf. Harman (1986).

[46] Both Boult (2020) and Schmidt (2021) discuss the possibility of subjecting an individual to epistemic blame in instances involving trivial truths. Importantly, however, the cases they focus on are very different from the case of Katlyn. The cases they focus on involve individuals who: (i) *believe* certain trivial claims which are (ii) *not* well supported by their evidence. The case of Katlyn involves a person who: (i) *fails* to believe certain trivial claims which (ii) *are* well supported by her evidence. I agree that there will be instances of the first kind where the person is appropriately subjected to a distinctly epistemic form of criticism. My claim is that *Katlyn* is not appropriately subjected to such a response so long as her tendency to avoid taking on beliefs that are well supported by her evidence is restricted to claims that are trivial and pointless.

[47] For a recent defense see: Piovarchy (2021).

features of our situations. In the epistemic case, the relevant blame-responses discourage certain forms of epistemic behavior and this, in turn, cultivates epistemic agency. If we understand epistemic blame in this way, Katlyn will, once again, not be appropriately subjected to this form of response. A "well-cultivated" epistemic agent is not one who will draw out the consequences of *any* piece of strong evidence she just so happens to possess, no matter how trivial or uninteresting. Katlyn can perform as she did while also being a fully competent and mature epistemic agent.[48]

Thus, I conclude that Katlyn is not appropriately subject to any form of criticism – even a *distinctly epistemic* form of criticism – on account of her not believing certain disjunctive claims. This is so even though Katlyn is being prompted with a certain question while also possessing (and consciously and vividly appreciating the evidential import of) strong evidence E for the truth of p[49]. This concludes the first step of my argument. We have now established the following claim:

> There are cases where: (i) a subject S possess strong evidence E for the truth of p at time t, (ii) all other minimal conditions for the normativity of E are met at t, (iii) S doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at time t.

3.2 An Argument Against Minimalism: Step 2

The second step of my argument will involve the following claim which links openness to criticism with epistemic reasons for belief:

> **If** [there are cases where: (i) a subject S possess strong evidence E for the truth of p at time t, (ii) all other minimal conditions for the normativity of E are met at t, (iii) S

[48] The agency-cultivation view of epistemic blame might be able to account for instances in which we blame a person for *believing* some trivial claim that's *not* well supported by the evidence; doing so perhaps reinforces a general tendency to respect our evidence. However, the case of Katlyn is importantly different from instances like this (see note 46 above).

[49] While I won't argue for it at length here, I believe that the considerations adduced in relation to the case of Katlyn and the pointless disjuncts apply to a wide range of instances involving so-called "trivial" truths.

> doesn't believe p at t, yet (iv) S isn't open to any form of criticism on account of (i)-(iii) at time t] **then**, in such cases, it's not the case that E is (or gives rise to) an epistemic reason for S to believe p at t.

While I think there is some plausibility to this linking claim, step 2 of the argument will not depend upon its truth[50]. Rather, I will argue that the minimalist can either accept or reject the linking claim. Either way, the minimalist view falters.

But why think there's *any* plausibility behind this linking claim in the first place? The driving idea behind this linking claim is that normative reasons for belief should somehow show up in our practices of interpersonal criticism and assessment. Their "normativity" should be revealed in the ways that we respond to individuals who either conform or fail to conform to them. If there is no trace of such reasons in our actual practices, then by what right can we call them "normative"? Consider, for instance, Schmidt (2021) who makes a similar point about moral normativity:

> the significance of a moral requirement will make it often – in absence of an excuse or exemption – appropriate to show resentment or indignation. These emotions are expressions of the normative significance we attach to the moral requirement because they are appropriate in face of its violation.[51]

Similarly, if *evidence* is normative (as the minimalist contends), then we should expect this to be somehow expressed within our interpersonal practices of criticism, blame, and accountability. When it comes to the epistemic (as opposed to the moral), such expressions don't necessarily have to involve certain *emotional* responses such as resentment and indignation[52]. Nevertheless, if evidence constitutes genuinely normative reasons for belief, then we should expect there to be

[50] As its stated, the linking claim is almost certainly too strong. I will return to this issue below.
[51] Schmidt (2021) pp. 4-5.
[52] For instance, on Boult's "relationship-based" approach to epistemic blame, the blamer doesn't have to be emotionally exercised.

some kind of expression of its normative significance within our practices. This is what's motivating the attempt to link openness to criticism with epistemic reasons for belief.

Thus, the above linking claim perhaps has some initial plausibility. Let us move on now to consider the two options mentioned above for the minimalist: either accepting or denying the linking claim. Recall that evidential minimalists subscribe to either one of the following:

> Strong evidence for the truth of p is *itself* an epistemic reason for S to believe p.
>
> Or
>
> Strong evidence for the truth of p *gives rise to* an epistemic reason for S to believe p when certain other minimal conditions are met (e.g. S possesses the evidence and considers whether p).

Consider a minimalist who accepts the linking claim. If the minimalist accepts the linking claim, and if the arguments presented in sec. 3.1 are on the right track, then she must abandon her position. This is because, as we have seen, *there are* cases where conditions (i)-(iv) are met. In other words, the antecedent of the linking claim is satisfied. However, this means that a minimalist who accepts the linking claim must infer the consequent, viz. that, in the relevant cases, it is *not* the case that evidence E is (or gives rise to) an epistemic reason for belief. However, this is incompatible with both of the claims in the disjunctive characterization of the minimalist position. Thus, it seems that the minimalist must *deny* the linking claim. I would like to consider three different routes the minimalist might pursue in denying the linking claim. As we will see, none of these routes will prove to be favorable for the minimalist.

First, the minimalist might reject the linking claim as too strong. Specifically, the minimalist could argue that there are cases where conditions (i)-(iv) in the antecedent are met, yet the relevant evidence *still* gives rise to an epistemic reason for belief. Imagine cases involving *excuses*. Say that a person is temporarily impaired, exhausted, or confused. Being in

such a state can lead one to violate certain standards or norms. However, this can also excuse a person from criticism. Imagine a friend who is suffering from depression or who is stricken with grief. Say that, as a result, the friend fails to follow through on some promise that was made to you. Perhaps a norm has been violated here (a norm that the person had good reason to follow), but the person is plausibly excused from criticism given their situation. Thus, the excuse doesn't imply that there *wasn't* a reason for the person to follow through on the promise, it just implies that the person *shouldn't be criticized* for failing to do so.

Similarly, there might be cases where conditions (i)-(iv) in the antecedent of the linking claim are met partly in virtue of the fact that the person in question is in a state which excuses them from criticism. The minimalist could argue that there are still epistemic reasons for belief in such cases (just as there was still a reason for the person in the above case to follow through on the promise). The point to emphasize here is that, as long as the minimalist still accepts some general connection between normative reasons for belief and openness to criticism, then appealing to excuses won't save her view. We could even introduce a fifth condition into the antecedent of the linking claim: (v) S is not excused or exempted from criticism.[53] Return again to the case of Katlyn and the disjunctive claims. This is a case where an agent is clearly and vividly attending to some evidence E for the truth of p without also believing p and yet is not appropriately subjected to criticism on account of this. Moreover, *there are no* excuses or exemptions at play in this example. For instance, we can say that Katlyn is an ordinary adult who is competent, intelligent, under no duress or stress, isn't impaired, etc. Cases like this show us

[53] Exemptions differ from excuses in being *global* rather than *local*; an exempted individual lacks the general capacities and abilities that are required in order to be appropriately subjected to ordinary interpersonal criticism and assessment. For instance, very young children and other animals would be *exempted* from criticism rather than *excused*. The conditions which give rise to excuses are, by contrast, usually temporary or "one-off".

that the appeal to excuses won't save a minimalist who endorses some general connection between normative reasons for belief and openness to criticism. In such cases, there are no excuses, yet the person *still* isn't appropriately subjected to criticism.[54]

Consider a second way a minimalist might try to reject the linking claim; a way which also regards the claim as too strong. A minimalist could reject the linking claim by saying that evidence only ever provides *warrant* for belief. For instance, in my (2021) I discuss the notion of a "warranting" reason[55]. A warranting reason to believe p doesn't *require* a person to believe p. Rather, warranting reasons are reasons that it would be appropriate to base one's doxastic attitudes on, whether or not one actually forms the attitude[56]. Arguably, Katlyn still has a warranting reason to believe the disjunctive propositions even if she doesn't take on beliefs in them, and even if she's not open to any form of criticism in virtue of this. Moreover, one could say that "warranting" reasons are a kind of normative reason for belief. Thus, according to this minimalist response, the linking claim should be rejected since the antecedent can be satisfied (along with our additional clause (v) ruling out exemptions and excuses) while the consequent is false. In the case of Katlyn and the disjunctive claims, the antecedent is satisfied, however the

---

[54] Why couldn't the minimalist just insist that there *are* excuses in cases like Katlyn's? Paradigmatic excuses include things like (non-culpable) ignorance and temporary impairment/lapses of judgment. As I've indicated, nothing like this obtains in the case of Katlyn. Thus, in order to maintain the claim that there *are* excuses in cases like Katlyn's, the minimalist would have to hold that *indifference* or *lack of interest* could function as an epistemic excuse. However, it should be clear that such a view is deeply at odds with evidential minimalism. The minimalist holds that evidence is normative for belief *regardless* of a person's idiosyncratic desires/interests/goals (recall again Kelly's points regarding MOVIE SPOILER mentioned earlier). Thus, the above appeal to excuses is unavailable to the minimalist.

[55] I borrow the term "warranting reason" from Abramson and Leite (2017).

[56] Note that warrant is stronger than mere permissibility. Something is permissible just in case it is not forbidden. Warranting reasons provide more than mere permissibility; they make *appropriate* or *justify* adoption of a certain attitude without *requiring* its adoption.

evidence *does* give rise to an epistemic reason to believe p, viz. a *warranting* reason.

I'm actually inclined to think that this response is correct. However, the point to emphasize here is that it won't help the minimalist. In order to see why, we need to ask the following question: under what conditions would a person be warranted in believing p? Return again to BORED AT HOME and LAST DONUT. In the former case we imagined Katlyn being asked "How many hours have you been alive?", in the latter case we imagined David being asked "Who ate the last donut?". In both cases, the individuals possess strong evidence E for the truth of p, however they must carry out some mental act M (e.g. a process of calculation or reasoning) in order to "see" the evidential connection between E and some target proposition p. The point I'd like to make here is that, in such cases, the person will only have warrant for believing the target proposition if they *actually engage* in the mental act M. Prior to engaging in that act, neither Katyln nor David would be warranted in believing their respective target propositions. However, engagement in such acts will often be a *non-minimal* condition. For instance, these will be acts that one might carry out when engaged in *active inquiry* vis-à-vis some question or subject matter. As we saw earlier, engagement in active inquiry is a non-minimal condition. Thus, there will be many cases where strong evidence E for the truth of p only gives rise to a warranting reason to believe p when some non-minimal condition is met, and this is incompatible with the minimalist's position.

Finally, a minimalist might reject the linking claim by simply denying *any* inherent connection between normative reasons for belief and our practices of interpersonal criticism, accountability, and blame. Consider, for instance, the objectivist described earlier; someone who thinks that strong evidence *just is* a normative reason for belief. This minimalist might hold that claims regarding our ordinary practices of interpersonal criticism and assessment are neither here

nor there when it comes to the nature of objective normative reality. If the minimalist takes this route, then the normativity of evidence becomes utterly mysterious. The forgoing discussion reveals that there are cases where conditions (i)-(iv) are met, the subject in question is consciously and vividly appreciating the evidential connection that obtains between E and p, and the subject is not exempted or excused from criticism. The case of Katlyn and the disjunctive claims is a case of this kind. The objectivist must persist in claiming that, nevertheless, in such cases the evidence E constitutes a normative reason for S to believe p. However, it's not clear what the objectivist could mean by "normative". If the above conditions are met, then in what sense does S have a "normative" reason to believe p?

The objectivist could appeal here to the notion of a warranting reason. She could, moreover, attempt to avoid the points I made above regarding warranting reasons by appealing to an *objective* notion of warrant. Thus, just as there might be objective "oughts", there might also be objective epistemic *warrants*. These would be (warranting) reasons to believe certain propositions which obtain for an agent *regardless* of whether or not the agent is properly apprised of them (e.g. by "possessing" them and appreciating their evidential import vis-à-vis some target proposition). However, this has the following unfavorable result: as I sit here typing, I have warranting reason to believe *infinitely* many propositions. Indeed, I have warranting reasons to believe things that I *can't,* given my psychological limitations, believe. Just imagine all of the pointless disjunctive claims that are entailed by the simple proposition "I am typing on my computer right now." The objectivist described above has to say that, as I sit here typing, I have warranting reason to believe *all of them*, and countless others. It's entirely unclear how the objectivist can maintain this while also taking "warrant" to be a normative notion.

4. Conclusion

In this chapter I have argued against a view that I have labeled "evidential minimalism". According to minimalists, there is a close connection between strong evidence for the truth of p and a normative reason to believe p: evidence is either *itself* is a normative reason for belief, or evidence gives rise to such a reason when certain other minimal conditions are met. My argument proceeded in two steps. In the first step, I established that there are cases where a subject S possesses strong evidence for p, all of the minimal conditions for the normativity of that evidence are met, yet S doesn't believe p and isn't open to any form of criticism on account of this. The second step involved a linking claim which connected openness to criticism with epistemic reasons for belief. I argued that the minimalist can either accept or reject the linking claim. Either way, the minimalist view falters. If my arguments against minimalism are successful, then some form of non-minimalism will be the way to go when it comes to understanding the normativity of evidence. In subsequent chapters, I will further explore the space of non-minimalist possibilities when it comes to understanding the normativity of evidence.

Chapter 2: Evidence and Epistemic Norms

1. Introduction

In the previous chapter I argued against a view that I labeled "evidential minimalism". According to evidential minimalists, there are either very *few* or even *no* conditions that need to be met, over and above the obtaining of strong evidence for the truth of p, in order for there to be an epistemic reason for S to believe p. The failure of the minimalist view means that evidence *itself* is not normative for belief. However, this is of course not to deny that evidence can *become* normative for doxastic agents. The task looking forward, then, will be to understand evidential normativity given the results of the last chapter.

Even if evidence is not itself normative for belief, there still might be correct norms of belief that incorporate the notion. Consider, for instance, the following candidate norms of belief:

(a) Don't believe without adequate evidence.

(b) Don't believe in a way that flies in the face of the total evidence in your possession.

(c) Don't reach a conclusion without having gone to appropriate lengths to seek out countervailing evidence.

All of (a) through (c) enjoin us, in some way or another, to be responsive to our evidence when forming and maintaining beliefs. Moreover, these norms seem importantly different from things like club rules or norms of etiquette. In some yet-to-be-clarified sense, they seem to have some kind of authority over us that isn't altogether easy to evade. However, if the minimalist view is mistaken, then it doesn't appear that their authority is traceable to something about *evidence itself*. Thus, we are left with the following question: if evidential minimalism is mistaken, then how are we to understand evidential norms of belief? The way I suggest moving forward with the discussion about evidential normativity is by considering which evidential norms of belief are *correct* norms of belief. Approaching matters in this way gives rise to a number of questions

related to norms and their correctness. The aim of the present chapter is to make explicit these questions and to outline how I intend to tackle them.

Before proceeding, however, I want to note something about the evidential norms above, (a)-(c). These evidential norms are formulated in a way that might lead one to think that evidential normativity can be understood by appealing to the nature of *belief*. For instance, (a) and (b) implore us not to believe in certain ways, and (c) concerns "reaching a conclusion"; something that we perhaps do when deliberating about what to believe. Thus, one might feel that the normativity of evidence is best understood by appealing to certain features of the attitude of belief itself. For instance, Bernard Williams famously claimed that "the characteristic of beliefs is that they aim at the truth."[57] A recent literature has been devoted to unpacking the idea that beliefs "aim" at the truth and exploring its various ramifications[58]. Depending on how this claim is unpacked, it might serve various explanatory ends, e.g. it could explain why we can't believe at will, or why evidential considerations take precedent in explicit doxastic deliberation.[59]

I do not, however, think that recourse to the aim of belief hypothesis will adequately address the problem of evidential normativity. The basic problem is this: evidence can sometimes give rise to normative claims concerning doxastic attitudes that we *don't yet hold*. A

[57] Williams (1973) p. 148.
[58] See, for instance: Shah (2003), Shah and Velleman (2005), Steglich-Petersen (2006), (2009), Velleman (2000), and Wedgwood (2002). See also the contributions in Chan (2013).
[59] The "aim of belief" approach is sometimes deployed in a way that is roughly analogous to certain Kantian approaches to moral normativity. According to a broadly Kantian approach to moral normativity, moral norms are constitutive of practical agency as such (see, for instance: Korsgaard (1996), (2009)). One could attempt to make a similar claim in the theoretical domain, viz. that certain norms or standards constitutively govern belief or cognition as such. For a recent defense of this view see: Horst (2022). There is an objection to this kind of view which I won't discuss in the main text (see: Enoch (2006), (2011)). The rough idea is that there seem to be many activities that are constitutively governed by certain standards (e.g. games) but which don't give rise to genuine *norms* or *reasons*.

person S could be "in the clear" vis-à-vis (a)-(c) by apportioning her current stock of beliefs to the evidence and by not reaching any further conclusions. Say some new evidence E comes to light; there is now conclusive evidence for the truth of p. Say also that S's existing beliefs do not conflict with this evidence; she doesn't hold views that require revision given the new evidence. According to (a)-(c) alone, then, there's no claim whatsoever upon S when it comes to believing p. This reveals a limitation of the "aim of belief" approach; it will fail to account for instances in which an agent has good reason to *expand* her current view[60]. In the last chapter I argued that there are *certain cases* where an agent possesses strong evidence E for the truth of p but there's no reason to believe p[61]. However, this is not to say that there *aren't* cases where evidence gives rise to a reason to take on a belief that we don't already hold. It seems, then, that evidential norms of belief cannot be understood solely on the basis of *either* evidence *or* belief. This points us in the direction of *some* form of non-minimalism, i.e. some view which ties the normativity of evidence to something beyond "purely" cognitive considerations.

In order to home in on the normativity of *evidence*, then, it will be better to focus on a norm that says something about how an agent should respond to strong evidence E for the truth of p even when there's no *revision* of the agent's current stock of beliefs that needs to be made in

[60] An aim of belief theorist could try to account for this by formulating a truth norm in a certain way, e.g. "For any p, one ought to believe p if and only if p is true". However, this is clearly too strong; it has the implication that we ought to believe *every* true proposition. The task for the aim of belief theorist would be to establish a norm that's capable of accommodating the idea that we sometimes have good reason to take on beliefs that we don't' yet hold – one that isn't *too* strong – and to do so in a way that somehow grounds the norm in the nature of belief itself. I won't consider at length whether or not the aim of belief theorist is capable of meeting this challenge. However, I am doubtful that any such norm could be successfully grounded in the nature of belief alone.

[61] For instance, prior to engaging in an explicit process of reasoning or calculation, both Katlyn in BORED AT HOME and David in LAST DONUT possess strong evidence E for some proposition p, yet neither of them has a reason to believe p.

response to E, only an *expansion* (something which (a)-(c) do not do). Keeping in mind the assumptions about evidence and evidence "possession" that were introduced in sec. 2 of the previous chapter, consider the following "Evidence Norm" (EN for short):

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

Note that, in many instances, in order for the evidential connection between E and p to be "clear to you" at time t, some *non-minimal* condition must be met at t. Recall again the case of Katlyn and the question of how many hours she's been alive. In order to "see" the evidential connection between E and p in this case, Katlyn had to engage in an explicit process of reasoning or calculation. As we have seen, engagement in such a mental act is a non-minimal condition. However, there will also be many cases where the evidential connection between E and p is clear to a subject S *without* the fulfillment of some non-minimal condition, e.g. instances where the connection between E and p is immediately apparent or obvious[62]. Recall also that a "warranting" reason to believe p doesn't *require* a person to believe p. Rather, warranting reasons are reasons that it would be appropriate to base one's doxastic attitudes on, whether or not one actually forms the attitude. Finally, "at the very least" is included in order to leave open the possibility of cases where a person *is* required to believe p when they possess strong evidence E for p and the evidential connection between E and p is clear.

EN, like (a)-(c), is an evidential norm of belief. However, it has a certain advantage over

[62] Imagine seeing some rabbit tracks in the snow. That there are tracks of such-and-so shape in the snow is strong evidence for the claim that there was a rabbit in the vicinity in the recent past. Someone who is familiar with the shape of rabbit tracks will, upon noticing the tracks, clearly "see" the connection between the evidence and the target proposition without much thought or effort. The evidential connection between E and p could also be made "clear to you" via the obtaining of some of the minimal conditions mentioned in the previous chapter, e.g. when someone prompts you with the question of "whether p" and you attend to E.

(a)-(c) in the present context; it says something about the conditions under which an agent has a reason to *expand* her current view. As a result, EN more clearly articulates the normative force of *evidence.* However, at this point, EN (as well as (a)-(c)) should be thought of as a *candidate* norm of belief; we have not yet established its correctness. The task going forward will be to determine whether or not EN *is* a correct norm of belief. As we have seen, it doesn't seem that evidential norms can be understood solely on the basis of *either* evidence *or* belief. Thus, if EN is a correct norm of belief, then some additional consideration(s) will have to be appealed to in order to understand its authority over us.

Pursuing matters in the above way gives rise to many questions concerning norms and their correctness. Sections 2-4 are devoted to addressing some of these questions. In section 2, I will relate evidential norms for belief to a broader class of *epistemic* norms. In section 3, I will distinguish two questions we can ask about epistemic norms: the "content" question and the "justificatory" question. In section 4, I'll note a difficulty confronting attempts to understand epistemic normativity and explain how I intend to overcome that difficulty. Finally, in section 5, I will consider, and reject, a recent attempt by Jane Friedman to identify the epistemic with the "zetetic", which pertains to inquiry.

2. Evidential Norm as Epistemic Norms

Pursuing the normative significance of evidence in terms of correct evidential norms for belief gives rise to a number of questions. Let's begin by considering the following "Club Norm" (CN):

CN: Gentlemen must wear sports jackets in the dining room.

We might identify this norm as a member of a particular class or category of norms, e.g. the norms governing the conduct of individuals who are members of a certain club. A club clearly

has limited scope, so the club's norms won't apply to non-members. However, an "outsider" could, qua anthropologist, come to have an understanding of the club and the norms that structure its operations. In other words, an individual S could come to understand what the norms of the club *are* without herself belonging to the club and thereby participating in its various activities. In such a case, S will not accept or internalize the club's norms. She will, nevertheless, understand what the norms are.

If we were to ask S whether or not CN is a "correct" norm there are a number of ways she might take this question. First, there's a fairly weak sense of "correct" which is simply descriptive; "correct" norms, in this sense, are whatever norms club members happen to accept. Second, there's a slightly stronger sense which we might call "correctness-as-permissibility". In this sense, even though S herself might not accept or internalize the norms, S acknowledges that it is permissible for others to do so. Finally, there's a still stronger sense of "correct" which is such that, if S were to acknowledge that the club's norms were "correct" in this sense, she would thereby be acknowledging that *she herself* has good reason to accept or internalize the norms.

If we're talking about the norms of a *club*, then whether or not S acknowledges their "correctness" in this last sense will depend upon various things about her, e.g. whether or not the club somehow appeals to her personal interests. To be clear, I want to leave open the possibility that S could be *mistaken* in her judgment that some norm or class of norms is "correct" in this last sense (as well as the first two senses). However, focusing now just on a good case (i.e. where S judges correctly), when it comes to the norms of something like a club, whether or not S acknowledges their "correctness" in the last sense described above will seem to depend upon things like S's personal interests. When it comes to something like a club, then, there will be many people who could correctly judge that its norms are *not* "correct" in the last sense

distinguished above.

Now consider the following "Promising Norm" (PN):

PN: If you promise S that you will Φ, then you ought to Φ.

If you promise someone that you will Φ but you fail to Φ, then you've violated PN. While CN might be considered a club norm, PN might be considered a "moral" norm, at least in the sense of that word which concerns what we owe to other people. Unlike CN, PN intuitively has a very *wide* scope of application. I've been availing myself of talk of norm "application", so allow me to explain what I mean by this. When I say that a norm N "applies" to an individual S, I mean that S can be *appropriately held accountable* for complying or failing to comply with N. "Holding accountable" is itself a very complicated subject matter; it comprises various forms of blame, sanction, and criticism[63]. Note as well that there will be instances where an individual S violates PN yet certain forms of response (e.g. blame) aren't called for, e.g. maybe S has some good *excuse*. Our practices here are varied and complex, but the basic idea is that PN "applies" to a wide group of individuals in the sense that there are many people who could be appropriately held accountable for either complying with or failing to comply with the norm.

When it comes to a norm like PN, it is more difficult to imagine an "outsider" who could come to recognize, qua anthropologist, the force of PN in the lives of a certain group of individuals while also correctly judging that she *herself* lacks reason to accept or internalize the norm. I won't argue that this is impossible. However, if it is possible, such individuals would seem to be quite rare. In other words, not only does a norm like PN appear to have a wide scope of application (at least when compared to a norm like CN), it also seems that there will be far

[63] Recall the discussion in the previous chapter involving several recent attempts to pinpoint a *distinctly epistemic* form of blame. These different proposals involved ways that a person could be held accountable for complying with epistemic norms.

fewer individuals – *if any* – who could correctly judge that PN is *not* "correct" in the third sense I distinguished above.

While the issues about the scope of application of some norm N and the range of individuals who could correctly judge that N is "correct" in the third sense distinguished above are separate matters, they also appear to be related in an important way. Return again to the individual described above who comes to have an understanding, qua anthropologist, of the norms that structure the activities of a certain club. Once again, let's say that this person could correctly judge that the club norms are *not* "correct" in the third sense described above. In this case, it does not seem appropriate to hold this person accountable for complying with the club's norms, i.e. the norms do not "apply" to her. Moreover, it seems that this is inappropriate precisely *because* the person has no good reason to accept or internalize the norms. In other words, there appears to be a connection between a norm's "application" to a person S and the question of whether or not S has good reason to accept or internalize the norm(s). We can formulate this "Application-Acceptance Connection" (AAC) as follows:

AAC: If a norm N applies to S, then there is good reason for S to accept N.

What AAC says is that, if a norm N "applies" to S (i.e. if it is appropriate for other people to hold S accountable for complying or failing to comply with N), then there is good reason for S to accept N. Notice that AAC does not say that, in order for a norm N to apply to S, S must *actually* accept N. Rather, AAC says that, in order for a norm N to apply to S, there must be *good reason* for S to accept N. There might be good reason for S to accept some norm N even if S does not, in fact, accept that norm. I'll say more about norm "acceptance" below. For now I'd simply like to note the following. If PN, unlike CN, has a wide scope of application, then there will also be far fewer people (at least when compared to CN) who could correctly judge that PN is *not* "correct"

in the third sense distinguished above.

When it comes to evidential norms for belief, it will also be helpful to have a way of classifying them as a family or a group (in the way that CN was classified as club norm and PN was classified as a moral norm). Presumably, evidential norms for belief are also different from those of a club, not only because the content of such norms doesn't pertain to things like dress codes, but also because, similar to moral norms, the scope of application of such norms seems to be much wider. In other words, it seems that evidential norms for belief will apply to a much wider class of individuals then the norms of a club[64]. Return again to the evidential norm (a) from earlier: Don't believe without adequate evidence. It seems that we can appropriately hold other people accountable for complying with this norm regardless of their personal goals, desires, and interests. If someone believes something without adequate evidence (say, that a certain vaccine causes some severe side effect, or that human induced climate change isn't real), then there are various forms of blame, protest, and criticism that are called for in response. Once again, the appropriateness of these responses doesn't seem to depend upon the target of the response possessing some idiosyncratic interest. Thus, evidential norms appear to apply to a wide class of individuals. Given AAC, what this means is that the class of individuals who could correctly judge that evidential norms are "correct" in the third sense described above will be a much larger class than, say, the class of individuals who could correctly judge that the norms of a particular club are "correct" in the third sense described above.

[64] I'd like to flag an issue here that I won't be able to return to until later (Chapter 5). One might think that talk of "accountability" in the epistemic domain is inappropriate given that epistemic evaluation concerns our beliefs and we lack voluntary control over what we believe. This is one further reason to take EN and (a)-(c) as *candidate* norms at this stage; one might hold that *there are no* epistemic "norms" (or at least none that *apply* to us) given that I'm connecting these notions to the concept of *accountability*.

We can do justice to this preliminary observation by saying that evidential norms of belief are a part of a wider class of *epistemic* norms of belief. According to a common construal, the "epistemic" pertains to certain cognitive achievements such as knowledge and understanding. As we will see below, recourse to the term "epistemic" won't automatically clarify our task since it's a term of art. Unlike, say, 'moral' or 'legal', the term isn't a part of our common parlance and thus can't be anchored in ordinary use. However, for the time being, we can make recourse to it as a way to capture the desired wideness of scope for evidential norms for belief. At least within philosophy, the "epistemic" is often thought to pertain to knowledge or intellection *as such*. To be clear, I fully grant that, upon reflection, "epistemic" norms of belief may end up having a scope of application that's much narrower than initially thought. Indeed, this is one of the tasks looking forward; to determine how wide or narrow the scope of epistemic norms actually *is*. However, "epistemic" at least gives us a label for the class of norms that evidential norms belong to, one that appears to have prima facie promise when it comes to securing a fairly wide application.

Thus, when trying to determine whether or not EN is a correct norm of belief, we will be trying to determine whether or not it is a correct *epistemic* norm of belief. EN should thus be treated as a member of a family of norms that are unified under the heading "epistemic". In the next section I will distinguish two different questions that we can ask about epistemic norms: the content question and the justificatory question. The aim is to get clearer on what "correctness" consists in when it comes to epistemic norms. We have, however, already seen one intuitive desideratum for epistemic norm "correctness": In order for some norm N to be a correct *epistemic* norm, N should "apply" to a fairly wide range of people. Given AAC, that means that there should be a fairly wide range of people who could correctly judge that N is "correct" in the

third sense distinguished above.

3. Two Questions Concerning Epistemic Norms

The first thing I'd like to do is separate two different questions we can ask about epistemic norms. While these questions aren't unrelated, it is important to distinguish them. First, there is a question about what the norms *are* or what they *say*. In other words, this question asks about the *content* of epistemic norms. Consider, for instance, two candidate epistemic norms, a truth norm and a knowledge norm:

TN: Don't believe p unless p is true.

KN: Don't believe p unless you know p.

In virtue of their incorporating factive notions like 'truth' and 'knowledge', TN and KN could be considered *externalist* norms of belief; in order to comply with them, our beliefs must meet certain objective criteria.[65]

It is, of course, controversial whether or not TN and KN actually *are* epistemic norms of belief. For instance, some think that these norms are too demanding. Consider, for instance, KEYS:

> KEYS: After coming home one evening, Jamie puts his keys down on the kitchen table and heads up to his bedroom. Shortly thereafter, Jamie's roommate takes Jamie's keys off the table, mistaking them for his own, and leaves the house. Jamie remains upstairs none the wiser, believing that his keys are still on the kitchen table.

In believing that his keys are on the kitchen table, Jamie violates both TN and KN. However, it doesn't seem appropriate to criticize him for this; he certainly isn't *blameworthy* for believing as he does. On the basis of examples like this, one might argue that TN and KN *aren't* norms of

[65] "Externalism" in epistemology is sometimes associated with the view that knowledge or justification do not require reflectively accessible grounds or reasons on the part of the subject. I don't mean to restrict myself to this conception of "externalism". For instance, compliance with KN might require reflectively accessible grounds or reasons (e.g. grounds or reasons that are somehow *objectively* adequate).

belief[66]. To be clear: I'm not claiming that such an argument would be a *good* one.[67]. However, setting certain worries to the side for the moment, we can at least see what the upshot of such an argument *would* be, were it successful: epistemic norms would be internalist in nature. In other words, norms like TN and KN *wouldn't* count as genuine epistemic norms. Instead, epistemic norms might require things like internal coherence or responsiveness to "one's" evidence, where this is construed internally.

Of course, the above argument might be mistaken and a broadly externalist view regarding epistemic norms might be correct. If that were so, then objective criteria would be included in the content of epistemic norms. My immediate aim is not to get ensnared in certain epistemological debates, but simply to point out *one kind* of question we can ask about epistemic norms, viz. a question about what they *are* or what they *say*. Notice that sorting out the internalism/externalism debate in epistemology wouldn't automatically settle this question. For instance, externalists might agree when it comes to the inclusion of objective criteria in the content of epistemic norms, but disagree over the finer details regarding the content of such norms (e.g. whether they *require* or merely *permit* certain beliefs). Thus, settling the internalism/externalism debate wouldn't automatically answer the content question about epistemic norms. However, as the above hopefully shows, the internalism/externalism debate is importantly relevant. This, then, is the first question I would like to isolate regarding epistemic

[66] A similar argument is known as the "New Evil Demon" argument (see: Cohen (1984)). The New Evil Demon argument is directed specifically towards reliabilist theories of epistemic justification. However, the underlying thought is similar: Given that blameless and responsible doxastic agents can run afoul of certain objective criteria in believing as they do, we should reject norms whose content incorporates such criteria.

[67] One way to push back would involve appealing to *excuses*. Being blameless for doing or thinking X doesn't mean that X was *the thing to be done or thought*; one could merely be *excused* for doing or thinking X. For discussion related to epistemic excuses see: Boult (2017), Greco (2019), Littlejohn (Forthcoming), Williamson (Forthcoming).

norms. Since it concerns the content of such norms, I will it the "content" question.

When it comes to EN, both internalists and externalists could be happy with the norm. This is because its content is hybrid: As per my assumptions in the previous chapter, I am operating with a broadly externalist view of evidence. However, EN also includes internalistic notions such as "possession" and a stipulation regarding the evidential connection's being "clear" to a subject. Thus, the norm is somewhat ecumenical as far as its content goes[68]. However, this does not mean that it's uncontroversial. Recall again the case of Katlyn and the pointless disjunctive claims discussed in the previous chapter (as well as other cases involving so-called "trivial" truths). One might hold that, in such cases, the evidence constitutes *no normative reason for belief at all* (not even a "warranting" reason). This is an option that is sometimes taken up by those who argue for broadly *instrumentalist* approaches to epistemic normativity[69]. Consider, for instance, Cowie (2014):

> Instrumentalists will of course admit that there is a sense in which 'evidence' and 'epistemic reasons for belief' can be used synonymously. They will deny, however, that this entails the existence of genuinely normative reasons to believe in accordance with one's evidence. Insofar as 'evidence' and 'epistemic reason for belief' are synonymous, there is reason to think that epistemic reasons for belief—as opposed to normative reasons for belief—are not necessarily normative.[70]

Thus, even though EN incorporates content that both internalists and externalists might be happy with, it is not uncontroversial. As the quote from Cowie indicates, instrumentalists might deny that strong evidence E for the truth of p is normative for belief in *any* sense, even when the

[68] Not only this, EN just specifies a *sufficient* condition for the existence of a normative (warranting) reason for belief.
[69] I'll say more about instrumentalism in the next chapter.
[70] Cowie (2014) p. 4014. This possibility is also discussed in Leite (2007) and Papineau (2013). Paakkunainen (2018) argues against the distinction between "genuinely normative" and "not genuinely normative" epistemic reasons for belief.

evidence is possessed and the connection between E and p is clear to the subject[71]. Additionally, in determining whether or not EN is a correct epistemic norm of belief, we will have to somehow secure its content. In carrying out this task, we will have to say something about determining the content of epistemic norms *in general*. Thus, carrying out this more general task will perhaps bring us more directly into contact with the internalist/externalism debate in epistemology[72].

There is another question concerning epistemic norms that I would like to isolate, one that is distinct from the content question. There are instances when a norm makes a claim on you, e.g. "Do X", and it's sensible for you to ask *why* you should do X in a way that calls into question the norm itself. This "why" question is the second kind of question regarding epistemic norms. Of course, norms don't simply say things like "Do X", they usually take something like the following form:

N: If condition C obtains, then do X.

There are cases where the kind of "why" question I have in mind, asked in response to a claim to do X, is not adequately answered by simply responding "Because C obtains". In such cases, what is being asked for is *a reason to accept N*. If there's no reason to accept N in the first place, then, even if C obtains, there won't be any good reason to perform X (assuming, of course, that there aren't other reasons to perform X which are unrelated to what's stated in N).

What is it to "accept" a norm? Having attitudes or performing actions that *conform* to the norm is neither necessary nor sufficient. Consider the following "Strong Evidence for Belief" (SEB) norm, where "strong evidence" should be understood in the way introduced in the last

[71] Indeed, in the next chapter, I will argue that this is what an instrumentalist *must* say; I will argue that instrumentalists do not have the resources to accommodate a norm like EN.
[72] I'll return to this issue in Chapter 6.

chapter[73]:

SEB: Don't believe p unless you possess strong evidence E for the truth of p.

One's belief that p might conform to SEB simply as a result of blind luck. For instance, one might be irresponsible in believing that p (and in one's believing more generally), basing it on either flimsy evidence or on no evidence at all. Even so, one might still possess strong evidence E for the truth of p. For instance, at the time of forming the belief that p, one may have simply failed to notice the connection between E and p. Thus, conformity to SEB is not sufficient for *acceptance* of SEB; if the person is a reckless and irresponsible believer (and let's say they're not entirely self-deceived about this) then, intuitively, they don't accept SEB. Conformity is also not *necessary* for acceptance of SEB. One might be a very responsible believer but hold a belief that p without possessing strong evidence for p; one might simply be non-culpably ignorant of the objective evidential situation vis-à-vis p. Return again to Jamie in KEYS. Since Jamie's belief is false, he doesn't possess "strong" evidence for it. In believing that his keys are on the kitchen table, he thus violates SEB. However, it doesn't follow from this that Jamie doesn't *accept* SEB. Perhaps Jamie is an extremely responsible and conscientious doxastic agent; he tries very hard to only believe things that are supported by strong evidence. Even if this is the case, there will be certain cases (e.g. KEYS) where Jamie does not conform to SEB.

Thus, "acceptance" of some norm N by S is some kind of attitudinal state had by S vis-à-vis N; one which does not guarantee S's *conformity* to N (try as we might, external circumstances aren't always congenial) and which is also not gotten simply *via* conformity to N (one can conform to a norm because of luck, accident, or coincidence). I won't try to further specify what this attitudinal state consists in here; I will simply assume that there is some such

[73] Recall that, in order for E to be "strong" evidence for the truth of p it must be the case that: (i) E is a true proposition (possibly about the external world), and (ii) one could come to *know* p on the basis of E.

state which counts as an individual's "acceptance" of a norm. I will also assume that there are norms which we have good reason to accept, and that there are norms which we do not have good reason to accept.

The "why" question that I have in mind asks about the *reasons* for accepting a norm or class of norms. Return again to the individual mentioned earlier who comes to achieve, qua anthropologist, an understanding of the norms that structure the activities of a certain club. Call this person S. S can meaningfully ask why *she* should bother to accept the club's norms. Thus, there is clearly a meaningful "why" question when it comes to S's acceptance of the club norms; if there's no good reason for S to join the club in the first place, then there won't be any good reason for her to accept the norms that structure its activities. Given AAC, this also implies that the norms will not "apply" to her. Consider, by way of illustration, the following "Meeting Norm" (MN):

MN: Meet at club headquarters every Thursday evening at 9:00 PM.

MN says that, if it's Thursday evening, then you should go to club headquarters for a meeting. Say that Thursday evening rolls around. S might explicitly acknowledge that it's Thursday evening but then *deny* that she ought to go to club headquarters. She might also be *correct* in denying this; if there's no good reason for her to accept the club's norms, then it's not the case that *she* ought to go to club headquarters, even if it's Thursday evening.

In other words, when it comes to the norms of a club, we can easily imagine scenarios where the following four conditions are met:

(i) There's some club norm N of the form "If condition C obtains, then do X".

(ii) S knows that N is a club norm.

(iii) S explicitly and correctly acknowledges that condition C obtains at time t.

(iv) S explicitly and correctly denies that *she* ought to X at t.

The example involving MN is an instance in which (i)-(iv) are met. The possibility of (iv), even given (i)-(iii), depends upon a certain negative answer to a question that S can ask vis-à-vis N, viz. a question about her reasons for *accepting* N. In other words, the possibility of (iv), even given (i)-(iii), depends upon there being no good reason for S to accept N in the first place. Even if this is so, S still might gain an understanding of what the norms *are* or what they *say* (this is captured in condition (ii)). Thus, in this case, it's quite clear that the content question is distinct from the kind of question that is currently under consideration. Since this latter question asks about the justificatory grounds (i.e. the reasons) for one's acceptance of a certain norm or class of norms, I will call it the "justificatory question".

While there might be a meaningful justificatory question when it comes to our acceptance of the norms of a *club*, one might doubt whether there's a meaningful justificatory question when it comes to our acceptance of *epistemic* norms. For instance, one could deny that there are parallel cases in the epistemic domain where the above conditions (i)-(iv) are met. If there are epistemic norms of the form "If condition C obtains, then do X", then the relevant conditions C might include things like possessing strong evidence E for the truth of p, and the relevant responses X might include things like believing p. However, one might insist that there's no "gap" that obtains between acknowledging that there's strong evidence E for the truth of p and acknowledging that one ought to believe p[74]. Such a "gap" clearly obtains between acknowledging that it's Thursday evening and acknowledging that one should go to club

[74] This issue connects with G.E. Moore's (1903) famous "open question" argument. According to Moore, ethical facts and concepts cannot be analyzed in terms of natural facts. For instance, for any natural fact N, it seems that we can coherently ask "X is good, but is it N?", i.e. this latter question is "open". There has been some discussion regarding open question arguments in epistemology. For instance, Greco (2015) argues that there are successful epistemic analogues of the open question argument. Heathwood (2009) and Côte-Bouchard (2017) disagree and argue that statements linking evidence/chance/probability and normative claims regarding what to believe give rise to "closed" questions.

headquarters. Moreover, one might appeal to the absence of this kind of "gap" in the epistemic domain in order to argue that *there is no* meaningful "justificatory" question when it comes to our acceptance of epistemic norms. In other words, one might argue that the obtaining of such a "gap" is a necessary condition for the existence of a meaningful justificatory question vis-à-vis our acceptance of epistemic norms. Given that it doesn't obtain, there is no meaningful justificatory question vis-à-vis our acceptance of epistemic norms.[75]

I'm not convinced that the obtaining of such a "gap" actually *is* a necessary condition for the existence of a meaningful justificatory question vis-à-vis our acceptance of epistemic norms. However, I'm happy to concede the point here. Even if it is a necessary condition, I think that I have already provided good reasons for thinking that it is met. As I have argued in the previous chapter, the view which I have labelled "evidential minimalism" is incorrect. According to minimalists, evidence *itself* is normative for belief. In other words, according to the minimalist, there's a very tight connection between strong evidence E for the truth of p and a normative reason to believe p. As I have argued, this view is mistaken. This provides us with strong reason to accept the sort of "gap" between evidence and normative claims concerning what we should believe mentioned above. In addition to this, there's a plausible psychological story that might be told regarding the connection between evidence and belief; one which also supports the sort of

[75] Why might one think that the obtaining of such a "gap" is a necessary condition for the existence of a meaningful justificatory question vis-à-vis our acceptance of epistemic norms? If there's some norm N of the form "If condition C obtains, then do X", and if there's no "gap" in between acknowledging that C obtains and acknowledging that one should X, then it seems that the only way that one could *avoid* acknowledging that one should conform to N is by avoiding acknowledging that condition C obtains. However, for creatures like us, there's no real possibility of avoiding the acknowledgment that there's strong evidence E for the truth of p; this kind of judgment is unavoidable. Thus, one could argue that, if the relevant kind of "gap" doesn't obtain in the case of epistemic norms, then there would be no way to avoid acknowledging that one should conform to epistemic norms; acceptance of epistemic norms would be inevitable and unavoidable. Perhaps "acceptance" wouldn't even be the right word to use in this context.

"gap" mentioned above. Perhaps we are psychologically constituted such that, when we confront (what we recognize to be) strong evidence E for the truth of p, we tend to form a belief that p. However, this psychological tendency isn't enough to secure a *normative* connection between evidence and belief. Just because we are psychologically disposed to believe in a certain way doesn't mean that we *should* believe in that way. For instance, perhaps we are psychologically constituted such that we are disposed to commit certain inferential fallacies. It of course does not follow from this that we *should* commit these fallacies.

Thus, there is a second question that we need to distinguish about epistemic norms; one that is distinct from the content question. As I have mentioned, I will call this question the justificatory question. When it comes to EN and the other evidential norms (a)-(c), then, it seems that we can meaningfully ask what reason we have to accept them, i.e. we can meaningfully ask a justificatory question about our acceptance of evidential norms of belief. Since evidential norms of belief make up a (perhaps large) part of the class of epistemic norms, it thus seems that we can meaningfully ask why we should accept epistemic norms in general, i.e. we can meaningfully ask a justificatory question about our acceptance of epistemic norms in general.

Let's take stock before moving forward. I began this chapter by noting that the problem of evidential normativity is best approached by considering which evidential norms of belief are *correct* norms of belief. I also highlighted a certain advantage that the following candidate norm EN has over other evidential norms (e.g. (a)-(c)):

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

The advantage is that EN more clearly articulates the normative force of *evidence.* The question, then, is whether or not EN is a correct norm of belief. I noted that it doesn't seem that evidential norms of belief can be understood solely on the basis of either evidence itself (as a minimalist

might argue) or on the basis of the nature of belief (as an "aim of belief" theorist might argue). Thus, we must appeal to some further consideration(s) in our attempt to understand the normativity of evidence. I then noted that evidential norms are best thought of as belonging to a wider class of *epistemic* norms. This led to a discussion regarding two different questions we can ask about epistemic norms: the content question and the justificatory question.

Looking forward, we will want to determine whether or not EN is a correct epistemic norm of belief. In order to determine this, we will want to say more about how to answer the content and justificatory questions vis-à-vis epistemic norms of belief. As I have noted, epistemic norms of belief appear to have a fairly wide scope of application. Given the "Application-Acceptance Connection" (AAC), this means that there will be a wide range of individuals who have good reason to accept epistemic norms of belief:

AAC: If a norm N applies to S, then there is good reason for S to accept N.

Thus, one key desideratum for a norm N's being a correct *epistemic* norm is that N be a norm that many people have good reason to accept. What this means is that, in order to establish EN as a correct *epistemic* norm, we need to develop a satisfactory answer to the justificatory question vis-à-vis EN for a fairly wide range of individuals. In other words, the reason(s) appealed to in our answer to the justificatory question vis-à-vis EN (and epistemic norms in general) must be shared by a wide range of individuals.

4. Beginning with the Justificatory Question vis-à-vis Epistemic Norms in General

Given the preceding, how should we proceed? One might think that we should start by tackling the content question vis-à-vis epistemic norms. After all, how could we even ask why we should accept a norm or class of norms when we don't yet know what such norms ask of us? Moreover, perhaps in answering the content question we will have *ipso facto* answered the justificatory question. While this might seem like a natural way to proceed, it will not be the

approach I adopt here. Instead, I'd like to begin with the justificatory question vis-à-vis epistemic norms for belief *in general* in the hopes that doing so might help us establish EN as a correct epistemic norm. I think there are good methodological reasons for approaching matters in this way, so allow me to say something about that.

First, the rationale offered above for starting with the content question actually misconstrues our pre-philosophical position. I take it as obvious that, in the context of ordinary life, we recognize certain norms as governing the doxastic lives of agents, e.g. prohibitions on bald inconsistencies, wishful thinking, or believing without good evidence. In other words, we commonly take violations of such norms as something that *calls for rectification*. Echoing a point I made earlier: we routinely *hold people accountable* for complying with norms like this. We could consider such norms as the "starting materials" for our investigation. This doesn't stack the deck in favor of certain philosophical views over others; these materials could be developed in various theoretical directions, and they are amenable to different conceptions of the "epistemic". Thus, our pre-philosophical starting position vis-à-vis the content of the relevant norms perhaps isn't quite as bleak as one might think. Second, my hope is that, by asking the justificatory question about epistemic norms first, we can perhaps ground epistemological debates regarding their content. If we can come up with a plausible answer to the question of why we should accept these norms, perhaps that will reveal something to us; it will reveal something about their proper place or purpose. This could then point us in the right direction when it comes to their content. My hope is that this might provide us with a way to move past mere appeals to intuition when it comes to philosophical discussions concerning the "epistemic" propriety of doxastic attitudes.

Thus, my hope is that, by starting with the justificatory question about epistemic norms for belief first, we will be assisted in our attempt to figure out whether or not EN is a correct

epistemic norm of belief. However, there's the following problem: there doesn't seem to be any non-arbitrary way of demarcating the "epistemic"[76]. In order to ask the justificatory question about epistemic norms in general, we need some kind of specification of the "epistemic" itself. Above I offered up a few commonsense norms that might be taken to be our "starting materials". However, as I mentioned, those materials could be developed in various theoretical directions, and they are amenable to *different* conceptions of the "epistemic". Thus, we do not seem to have, at the outset, some non-arbitrary way of understanding what unifies such norms as *epistemic* norms. Earlier I noted that, according to a common conception, the "epistemic" is linked with certain cognitive achievements such as knowledge or understanding. However, this might beg the question against certain internalist views which understand subjectivist notions such as rationality or coherence to be the distinctive purview of the epistemic. If at all possible, I'd like to avoid beginning with a stipulative conception of the "epistemic" which rules out certain approaches to the content question vis-à-vis epistemic norms in general (e.g. internalist approaches).

This issue constitutes a genuine roadblock for attempts to understand epistemic normativity. However, my view is that it can be overcome. In other words, I think there is a principled way to demarcate "epistemic" norms from other, non-epistemic, norms. I won't be in a position to fully defend my positive view until later. However, my idea is roughly this: Given the "starting materials" mentioned above, and beginning from a position within our everyday practices, we can ask ourselves the following: What kinds of normative expectations do we *hold other people to* when it comes to their doxastic lives? In the previous chapter I noted that there are various ways in which we engage in "criticism" of each other when it comes to the things that we believe, where by "criticism" I had in mind something more than mere assessment vis-à-

[76] For discussion see: Cohen (2016).

vis a standard. For instance, it seems that violating a norm of belief can sometimes appropriately give rise to certain *responses and reactions* on the part of others, e.g. various forms of blame, reproach, or sanction[77]. Moreover, it seems that certain classes of norms have *distinctive kinds* of reactions that are called for when they are violated, e.g. blame in the case of moral norms, or sanctions in the case of legal norms. My thought is that, if we can locate a distinctive kind of *response* that can be appropriately taken up in reaction to violations of norms similar to the ones I offered above as our "starting materials", then that will help demarcate the class of epistemic norms itself[78].

Approaching matters in this way could then help us answer the justificatory question about epistemic norms for belief in general. If we partake in a social practice wherein we respond to the norm violations of others in a distinctively epistemic manner, *and* the social practice is a legitimate one, then we will have an answer to the justificatory question in the offing: We should accept epistemic norms given that such norms structure a legitimate social practice that we participate in. It could also help us with the content question vis-à-vis epistemic norms. For instance, we can ask ourselves: What kinds of norms make the most sense of our actual interpersonal practices of epistemic criticism? This could then provide us with a way to fill out the content of the relevant norms.

I am thus of the opinion that philosophers of epistemic normativity would do well to take a cue from Peter Strawson who, in "Freedom and Resentment", wrote about certain "commonplaces" pertaining to our interpersonal lives. Strawson was concerned to point out the

---

[77] Also included in the relevant kinds of responses and reactions here are *excuses.* Excusing an individual for some norm violation thus does not mean the person isn't being *held to* some normative expectation. Additionally, there are various ways in which we can hold *ourselves* to various norms and standards (e.g. through guilt, self-reproach, etc.)

[78] This idea is also pursued by Kauppinen (2018). My approach will differ from Kauppinen's in certain respects which I will return to later.

significance of certain reactive sentiments such as resentment and indignation. Of these commonplaces he wrote:

> The object of these commonplaces is to try to keep before our minds something it is easy to forget when we are engaged in philosophy, especially in our cool, contemporary style, viz. what it is actually like to be involved in ordinary interpersonal relationships, ranging from the most intimate to the most casual.[79]

I do not mean to suggest that the kinds of reactions and responses I have in mind when it comes to *epistemic* norms are things like resentment and indignation. However, I *do* think that interpersonal responses will play an important role in answering both the justificatory and the content questions.

This, then, will be the way that I propose overcoming the obstacle mentioned above. However, before pursuing my approach, I would like to consider some possible alternative strategies for understanding the "epistemic". Considering how these alternative strategies fall short will help pave the way for my preferred approach to epistemic normativity. In the next section I will consider a view recently put forward by Jane Friedman in which she proposes identifying the epistemic with the *zetetic* (which pertains to inquiry). In the next chapter I will consider various way in which one might try to understand epistemic normativity as *instrumental* normativity[80]. I will argue against all of these competing strategies.

5. The Epistemic and Zetetic

In her (2020) Jane Friedman argues for a close connection between the epistemic and the "zetetic"[81]. The zetetic pertains to *inquiry*. Earlier I noted that "epistemic" norms at least appear to apply to a wide range of individuals. Friedman notes that human beings commonly engage in

[79] "Freedom and Resentment" p. 4.
[80] Friedman's view could also be considered an instrumentalist view. However, I would like to keep her view separate from the views I'll consider in chapters 3.
[81] Friedman has developed her zetetic approach to epistemic normativity in a number of recent papers: Friedman (Forthcoming), (2020), and (2019).

inquiry. We not only search for answers to the questions of philosophy and science; we also search for answers to routine matters such as the location of our keys and the time of day. Zetetic *norms*, then, are those norms that are binding on us *qua inquirers.*

Friedman highlights a tension between zetetic norms and what many contemporary epistemologists consider to be epistemic norms. Think about a typical case of inquiry where one wants or needs to figure something out. Consider Friedman's example:

> Say, for instance, that I want to know how many windows the Chrysler Building in Manhattan has (say I'm in the window business). I decide that the best way to figure this out is to head down there myself and do a count. To do my counting, I set up outside of Grand Central Station. Say it takes me an hour of focused work to get the count done and figure out how many windows that building has.[82]

Now consider the hour during which Friedman is doing her counting. During that hour there are many other ways that Friedman could make epistemic gains. She could, for instance, draw out consequences from things that she already knows in order to arrive at new knowledge. Also, the amount of perceptual information available to her at Grand Central during that hour is vast. Thus, there is a lot of new knowledge she could acquire and evidence that she could follow during her hour at Grand Central. However, during this hour, she doesn't do this. She doesn't draw inferences that aren't relevant to her counting task, and she does her best to ignore other things going on around her. Moreover, this seems like what she *should* be doing given the inquiry she's engaged in. If this is right, however, then during this hour she should be *ignoring* available evidence and information. Success at inquiry thus requires "behaving in something of an epistemically dubious way"[83].

The tension that Friedman identifies is thus between zetetic norms and norms that are commonly taken to be epistemic norms by contemporary epistemologists. For instance, here's

[82] Friedman (2020) p. 502.
[83] Ibid. p. 503.

her "zetetic instrumental principle" (ZIP):

> ZIP: If one wants to figure out $Q^?$, then one ought to take the necessary means to figuring out $Q^?$.

As the above example illustrates, complying with ZIP often requires us to *not* follow evidence that is easily available. ZIP is thus in tension with certain norms that are commonly offered as instances of "epistemic" norms. For instance:

> EP_a: If one has excellent evidence for p at t, then one is permitted to judge p at t.
>
> KP_a: If one is in a position to come to know p at t, then one is permitted to come to know p at t.

The subscript 'a' stands for 'act', indicating that these are norms for forming beliefs (which Friedman calls "judging") and coming to know. It looks as though, according to EP_a and KP_a, it would be perfectly permissible for Friedman to come to know anything at all that was available to her during her hour outside Grand Central, and perfectly permissible for her to follow any excellent evidence she had during that time. However, according to ZIP, during this hour such things are *not* permissible; the satisfaction of her inquiry-theoretic ends depends upon her *ignoring* much of her available evidence and *not* coming to know many things that she's in a position to know. Thus, if ZIP, EP_a, and KP_a are all *epistemic* norms, then there seems to be a tension within epistemic normativity.

How should we resolve this tension? Friedman thinks that separating the epistemic and the zetetic is not a good option. Instead of sequestering the two, Friedman argues for the *unity* of the epistemic and the zetetic. However, if one wants to argue for a unity between the epistemic and the zetetic *and* resolve the above tension, then it seems like either ZIP or norms like EP_a and KP_a will have to be rejected. Friedman supports a revisionist route that rejects norms like EP_a and

KP_a[84]. Denying norms like this amounts to denying that we are "always and everywhere epistemically permitted to judge p when the evidence clearly supports p, and that we are always and everywhere epistemically permitted to come to know."[85]

If we were to follow Friedman and unify the epistemic with the zetetic, how would that affect our understanding of epistemic norms of belief (including evidential norms of belief)? It seems that, if epistemic norms are zetetic norms, then we would only have reason to follow such norms *when we are engaged in inquiry.* Friedman introduces her discussion by noting that we are subjects "in pursuit" of information, and that we "want to know" things like where our keys are. Thus, inquiry vis-à-vis some question or subject matter is an intentional activity that one engages in; it is a purposive or goal-directed activity whose aim is the acquisition of certain information or the settling of a certain question. Indeed, in other work, this is how Friedman explicitly characterizes inquiry: "A true inquirer then is somebody with a certain kind of *goal or aim*, and so at the bottom of any true inquiry is a certain kind of *aim- or goal-directed state of mind or attitude*."[86] Given that our inquiries are circumscribed, and that our evidence base at any given time will include much that is irrelevant to any particular inquiry we might engage in, there will inevitably be strong evidence that we possess which does not bear on some inquiry that we're engaged in.

The problem with this is that there are many cases in which we place normative expectations on agents to respond to evidence for the truth of p *even when they're not engaged in inquiry* vis-à-vis p. In other words, it is a feature of our ordinary practices that

[84] Friedman states that this is revisionary vis-à-vis popular views in contemporary epistemology. As I will make clear below, I think it's also revisionary of our ordinary epistemic practices.
[85] Friedman (2020). p. 530.
[86] Friedman (2019) p. 298. (emphases added).

epistemic norms "apply" to agents even when they are not actively engaged in inquiry[87]. We already saw one example in the previous chapter. Recall MOVIE SPOILER:

> MOVIE SPOILER: Liz often sees newly released movies only after they've been in theaters for some time. Prior to her seeing some particular movie, she has no interest in forming true beliefs about its ending. In fact, she actively does *not* want to take on such beliefs. Liz is chatting with a few friends before class. One of them quickly changes topic and begins talking about a movie that Liz hasn't yet seen. Before Liz has a chance to warn her not to spoil the ending the friend blurts out the ending and Liz comes to form a true belief about its ending.

If we knew that Liz clearly heard her friend divulge the ending of the movie, then we would expect her to have a view on how the movie ends. This is so even though Liz wasn't actively engaged in inquiry regarding the matter. What's more, Liz actively *did not* want to find out about the ending of the movie. Still, it seems appropriate to expect her to have a view about the movie's ending if she clearly heard her friend[88]. Consider one further case:

> ROOMMATE: Kyle is a grad student. Kyle shares a house with one other grad student. During summer break, Kyle's roommate takes a road trip. Kyle's roommate is unsure how long his trip will take but estimates around two weeks. It's the middle of the second week and Kyle has gotten used to his roommate's absence and quite enjoys having the place to himself. While walking home one day deeply immersed in thought about a paper he's working on, Kyle all of a sudden spots his roommate's car in the driveway from down the street. He immediately comes to believe that his roommate is back in town.

In ROOMMATE, Kyle is not engaged in inquiry regarding the whereabouts of his roommate;

---

[87] Recall that, in order for a norm N to "apply" to S, it has to be appropriate for other people to *hold S accountable* for complying or failing to comply with N. I take it that holding a person accountable involves holding that person to a *normative* expectation. Such expectations are importantly different from non-normative expectations (e.g. expecting our car to make the trip, expecting our thermometer to give an accurate reading, etc.). I will say more in Chapter 4.

[88] Why think that this is a "normative" expectation? I won't be able to fully address this issue until later (Chapter 4). However, return again to the proposed accounts of "epistemic blame" from the previous chapter. If Liz avoids taking on beliefs that are strongly supported by her evidence simply because such beliefs would flout her personal desires, then it seems that she's believing in a way that opens her up to certain forms of response on the part of other people. For instance, perhaps other people would be justified in lowering trust in her when it comes to certain topics or subject matters. I will argue in Chapter 4 that this kind of response goes beyond mere assessment vis-à-vis a standard.

he's totally immersed in thought and not attending to the matter. Nevertheless, upon spotting the car, it's seems reasonable to expect Kyle to believe that his roommate is back in town. For instance, if we knew that Kyle spotted the car, then we'd expect him to have a certain opinion on the whereabouts of his roommate regardless if he was actively engaged in inquiry regarding this subject matter when he saw the car.

I offer these as some intuitive reactions to the cases offered above. What I think these cases suggest is that we hold people to certain expectations when it comes to responding to their evidence even when they are not engaged in an inquiry of a certain kind. In ROOMMATE, Kyle is perhaps engaged in inquiry (he's thinking about a paper he's working on), but the relevant evidence (the car in the driveway) doesn't bear on that particular inquiry in any way. In MOVIE SPOILER, Liz was not only *not* engaged in inquiry vis-à-vis the ending of the movie, she actively *did not* want to find out about the ending of the movie. Nevertheless, in both of these cases it still seems reasonable to expect the subjects to respond appropriately to the evidence. In other words, it still seems that epistemic norms "apply" to them. Given AAC, this means that the subjects have good reason to accept epistemic norms *even when* they are not explicitly engaged in inquiry.

As I said, these reactions strike me as intuitive. Later I will offer a more systematic way to think about certain of our reactions to doxastic agents. However, for now, I'll leave things at the intuitive level. The general thought is this: Our expectations of others when it comes to responding to evidence don't seem to sit well with a view that collapses the epistemic and the zetetic. Friedman is, of course, aware of this. As she mentions, if we reject norms like EP_a and KP_a, then "we will have to say that there may well be cases in which following our excellent evidence and coming to know will have been a mistake—a thoroughly epistemic mistake."[89]

[89] Friedman (2020) p. 530.

This appears to be a bullet she's willing to bite. While I'm not staunchly opposed to revisionism, Friedman's view seems to call for a quite radical revision; one which might call into question the legitimacy of at least some of the responses I've been considering above. My own view is that Friedman's position calls for too much revision and that there's an alternative conception of the "epistemic" which is preferable. However, I won't be in a position to fully defend my alternative view until later. As I said, I offer the above as an intuitive case against collapsing the epistemic and the zetetic.

Thus, if we understood epistemic norms as *zetetic* norms then it seems that epistemic norms will only apply to individuals who are engaged in a certain kind of purposive or goal-directed activity. This, however, appears to fly in the face of our ordinary expectations of one other when it comes to responsiveness to the evidence. Moreover, it also seems as though collapsing the epistemic and the zetetic is generally fraught with many difficulties. In the next chapter, I will continue the discussion and consider one further way we might try to understand the "epistemic". This will be the *instrumentalist* view. According to the instrumentalist view, our acceptance of epistemic norms is instrumentally justified. In other words, we have reason to accept such norms because doing so is an effective way to satisfy our aims, interest, or goals. This view shares some similarities with Friedman's approach insofar as it involves aims and goals. However, the view is more expansive than Friedman's insofar as it does not restrict itself to the particular aims or goals involved in the activity of inquiry.

6. Conclusion

I began this chapter by suggesting that the problem of evidential normativity is best pursued by considering which evidential norms of belief are correct norms of belief. I also organized the discussion around the following evidential norm EN:

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

The question then became: Is EN a correct norms of belief? I then argued that evidential norms of belief are a part of a broader class of *epistemic* norms. This led to a discussion about epistemic norms in general. I distinguished two questions about epistemic norms: the "content" question and the "justificatory" question. Thus, when it comes to determining whether or not EN is a correct epistemic norm, we need to keep both of these questions in mind. I then suggested that we should begin with the justificatory question vis-a-vis epistemic norms in general. My hope is that this will help us answer the content question for epistemic norms in general, and evidential norms in particular. However, I noted a roadblock confronting this approach having to do with the very idea of the "epistemic"; there doesn't appear to be a principled way of demarcating the class of *epistemic* norms. I gestured at my proposed solution to this problem, noting that I won't be in a position to fully develop it until later. According to my proposed solution, we can demarcate the epistemic by trying to locate a distinctive kind of response that can be taken up on the part of others in reaction to an individual's violation of a norm of the relevant kind. I suggested that this might help us with both the content question and the justificatory question vis-à-vis epistemic norms. Before developing my own account, I proposed considering alternative ways of construing the epistemic. I considered and rejected Friedman's attempt to unify the epistemic with the zetetic.

Chapter 3: Against Epistemic Instrumentalism[90]

1. Introduction

In the previous chapter I argued that the problem of evidential normativity is best pursued by considering which evidential norms of belief are correct epistemic norms. I then proposed beginning with the justificatory question vis-à-vis epistemic norms in general. I also considered, and rejected, Friedman's proposal which would have us collapse the epistemic with the zetetic. Before moving on to develop my own approach, I would like to explore the possibility of understanding epistemic norms as *instrumental* norms. One might grant that Friedman's view goes too far in collapsing the epistemic and the zetetic yet nevertheless feel that *something* in the general direction of her approach is on the right track. Recall that the "zetetic" pertains to inquiry. Our inquiry-theoretic ends are satisfied when we acquire certain information or answer certain questions. When Friedman figures out how many windows the Chrysler Building has, she will have satisfied her inquiry-theoretic end of answering that particular question. Friedman's proposal is that epistemic norms simply *are* zetetic norms, i.e. norms that are binding on us qua inquirers. In the last chapter I raised some difficulties for this view. The basic idea was that there are many cases (e.g. ROOMMATE and MOVIE SPOLIER) where we seem to expect agents to respond to the evidence *even when* they are not engaged in the activity of inquiry.

One could accept these points yet nevertheless persist in thinking that epistemic norms are still *instrumental* norms. As we have seen, Friedman's account is fairly restrictive insofar as it only incorporates the aims or goals associated with a particular kind of task or activity, viz. inquiry. However, one could appeal to a broader range of aims, interests, and goals in an attempt to bypass the worries that I raised in the last chapter. As we will see, this will be the strategy

[90] The following chapter draws heavily from my (2021). I would like to thank two anonymous reviewers at the journal *Synthese* for their helpful feedback and suggestions.

adopted by instrumentalists.

According to instrumentalists, our acceptance of epistemic norms is instrumentally justified; we should accept such norms given their usefulness when it comes to the satisfaction of our aims, interests, and goals. "Aims, interests, and goals" comprise a diverse set of attitudes that are sometimes called "pro-attitudes". A paradigmatic case would involve a *desire*. If I desire to learn a certain piece on the piano, and a particular practice schedule (e.g. one that prescribes a certain number of practice hours per week) represents an effective means via which that desire could be satisfied, then I would thereby be justified in accepting that practice schedule. Similarly, according to instrumentalists, we should accept epistemic norms of belief since they represent an effective way to satisfy our aims, interests, and goals (i.e. our "pro-attitudes"). While desire is a paradigmatic pro-attitude, the class is significantly varied; it encompasses preferences, wishes, intentions, inclinations, and more. I'll often use "aims, interests, and goals" as a convenient label for this varied class. According to instrumentalists, then, epistemic norms are binding on us given their instrumental efficacy in relation to the satisfaction of the attitudes that comprise this varied class.

A number of authors have recently defended versions of the instrumentalist view in epistemology[91]. The view comes in many different varieties. I will consider a few of those varieties in the present chapter. However, I don't purport to provide an exhaustive survey of the space of instrumentalist possibilities. My aim is to cover a few main versions of the view. However, before turning to the different versions of the view that I will consider, I would like to clarify how I'll be understanding the instrumentalist view. First, I will understand the "aims,

[91] Côté-Bouchard (2015), Cowie (2014), Dyke (2020), Foley (1987), Kornblith (2002), Laudan (1990), Leite (2007), Papineau (2003), Sharadin (2018), Steglich-Petersen (2018). Quine is also sometimes read as an epistemic instrumentalist (see especially his (1986)). For arguments against the instrumentalist view see: Buckley (2021), Kelly (2003), (2007), Paakunainen (2018).

interests, and goals" appealed to by the instrumentalist as (i) *personal level* aims, interests, and goals that are (ii) *had by individual agents.*

Regarding (i), one could attribute "aims" or "goals" to sub-personal states or processes. So, for instance, we might say that the "aim" or "goal" of the heart is to pump blood. Here, "aim" talk should be unpacked as proper functioning, or perhaps as correctness qua type of object in question; hearts don't literally "aim" at things in the way that archers do. As we saw in the last chapter, it is sometimes said that truth is the "aim" of belief. One way to understand the dictum that beliefs "aim" at truth is to say that beliefs have as their constitutive standard of correctness a norm which privileges the truth. The exact formulation of this norm is a matter of dispute and need not concern us here. The point to keep in mind is that, if beliefs do indeed have a truth norm as a constitutive standard of correctness, then this norm governs the beliefs had by an agent regardless of *that agent's* adopted aims, goals, or interests. Indeed, the norm would govern the agent's beliefs even if the agent had no interest in the truth. Thus, one could attempt to argue that we should accept epistemic norms by appealing to the constitutive norm of belief. Such a position would not, however, count as a version of instrumentalism, as I'm understanding the view[92]. As I'll understand the view, "aims, interest, and goals" occur at the personal level;

[92] There is another way of unpacking the claim that beliefs "aim" at truth which says that, insofar as an agent is *deliberating* about whether or not to believe p, that agent has the intentional, personal-level goal of believing p if and only if p is true. In other words, the activity of conscious belief formation has this truth-oriented aim as its "constitutive" aim or goal. Given that the goal in question is an intentional, personal-level goal had by an agent, I would count such a view as a version of instrumentalism. However, such a view seems ill-equipped to account for the normativity of evidence. First, most of our beliefs are not formed via explicit deliberation. Nevertheless, evidence can sometimes constitute reasons for or against such beliefs. Second, it seems that evidence can constitute reasons for us to take on beliefs that we don't already possess. Such beliefs may pertain to issues that we haven't deliberated about and do not plan to deliberate about. For a criticism of various attempts to utilize "constitutivist" strategies to explain epistemic normativity see: Côté-Bouchard (2016). For a criticism of a defense of epistemic instrumentalism that makes appeal to the personal-level version of the "aim of belief" thesis see: Côté-Bouchard (2015, §4).

they are the kinds of things that we could attribute to individuals utilizing commonsense folk-psychological discourse (e.g. "He wants X", "She desires that Y", etc.)[93].

Regarding (ii), the "aims, interests, and goals" invoked by the instrumentalist will also be had by *individual agents.* Thus, I will not engage directly with the instrumentalist view put forward in Dyke (2020) which grounds epistemic normativity in the aims or goals had by *collectives* or *communities* of individual agents. Thus, we could call the form of epistemic instrumentalism that I will be concerned with here *individualist* epistemic instrumentalism rather than *collectivist* epistemic instrumentalism.

I will also, by and large, be focused on instrumentalist views that appeal to the aims, interests, and goals that individuals *actually* have, rather than the aims, interests, and goals that they *should* have or that they *could* have[94]. For instance, one could ask what norms we would be justified in accepting if we *were* to have certain aims, goals, and interests. If individuals were to have the goal of believing all and only true propositions, what norms would they be justified in accepting? This is a fine question, but it's not one that I see the instrumentalist as trying to answer. Instrumentalists, as I will understand them, are trying to justify acceptance of epistemic norms by appealing to the aims, interests, and goals that individuals *do* have. Similarly, I will not focus on views which appeal to aims, interests, and goals that individuals *should* have. Again, I see the instrumentalist as trying to ground epistemic normativity in the aims, interests, and goals that individuals do, in fact, have. Indeed, this is often offered as one of the main selling points of

[93] I want to leave open the possibility that such ascriptions could be true even in cases where the target of the ascription isn't conscious of their own mental state. In such cases, the relevant want, desire, wish, etc. could still be "theirs" even though unconscious. The mental state would still occur at the "personal level" in virtue of the fact that it is the kind of state that, *in principle,* one could consciously hold.

[94] There's a slight caveat here that I'll discuss in section 3. The basic idea is that one might introduce a counterfactual element in a way that's amenable to the broad approach favored by the instrumentalist. I'll say more below.

epistemic instrumentalism: It represents a naturalistically friendly way of accommodating epistemic normativity. According to the instrumentalist, epistemic normativity is simply the normativity of means-ends efficiency, where the "ends" are provided by contingent facts about human psychology.

I would also like to note a way that one could argue against the instrumentalist. While this isn't a line of argument that I will pursue at length here, it does represent a serious challenge for the instrumentalist. The argument is that aims, interests, and goals can't, in and of themselves, *justify* anything. Take, for instance, the case of desire. Say that S desires to Φ and pursuing means M would be an effective way to satisfy that desire. One could argue that this alone would not justify S in pursuing means M. One could argue that whether or not S is justified in doing this depends crucially on the content of her desire. If her desire is to, say, rob a house and breaking a window would be the most effective way to achieve that goal, then that wouldn't *justify* S's breaking of the window; it might *rationalize* it, but rationalization and justification are two different things. Similarly, one could object to the instrumentalist's attempt to justify our acceptance of epistemic norms by appealing to our aims, interests, and goals; appealing to such attitudes would perhaps *rationalize* our acceptance of such norms, but it wouldn't *justify* our acceptance of them. I think this is a serious challenge for the instrumentalist to consider, but it's not one that I'll pursue here. In other words, I'm willing to grant the instrumentalist the justificatory relevance of aims, interests, and goals. My view is that, *even granting this*, the instrumentalist view falters.

This, then, is how I will understand the instrumentalist view. Recall that my overarching task is to try to understand evidential normativity. In the previous chapter I argued that this task is best approached by considering which evidential norms of belief are correct epistemic norms of belief. In particular, I focused on the following norm:

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

The question, then, is whether or not EN is a correct epistemic norm of belief. My suggestion is that we begin with the justificatory question vis-à-vis epistemic norms in general first. My hope is that, by beginning with the justificatory question about epistemic norms in general first, we will be assisted with the content and justificatory questions vis-à-vis evidential norms of belief (including EN).

The instrumentalist view offers us a possible way of understanding the "epistemic" as well as a potentially attractive answer to the justificatory question vis-à-vis epistemic norms in general. The basic instrumentalist suggestion is that "epistemic" norms of belief are those norms that we have good reason to accept given their usefulness when it comes to the satisfaction of our aims, interests, and goals[95]. According to the instrumentalist, this is why we have good reason to accept such norms; they represent an effective way to satisfy our aims, interests, and goals. As I mentioned above, I'm willing to grant to the instrumentalist the justificatory relevance of aims, interests, and goals. Thus, I won't pursue an argument against the instrumentalist which holds that the view can't answer the justificatory question because of the justificatory irrelevance of aims, interests, and goals.

How, then, to argue against the instrumentalist? The way I will argue against the instrumentalist is by showing that none of the main varieties of instrumentalism is capable of accommodating EN. In particular, for each variety of instrumentalism that I will consider, I

[95] To clarify: *Acceptance* of some norm(s) might not itself further our aims, interests, and goals given that an "acceptance" is simply some attitudinal state that one has vis-à-vis a norm. The idea is that accepting epistemic norms would likely have certain downstream effects on one's doxastic life (e.g. increasing the chances of acquiring true beliefs and avoiding false beliefs) and that these effects would help promote or satisfy our aims, interests, and goals.

will show that it cannot establish that a sufficiently wide range of people have good reason to accept EN. Recall that this was a key desideratum for epistemic norm correctness. As I have noted, epistemic norms appear to apply to a wide range of individuals. Given AAC, this means that, in order for some norm N to be a correct *epistemic* norm, a sufficiently wide range of people must have good reason to accept N. I will argue that none of the main varieties of instrumentalism has the resources to satisfy this desideratum when it comes to EN. In order to successfully answer the justificatory question vis-à-vis EN, the instrumentalist must show that a sufficiently wide range of individuals possess aims, interests, and goals that could underwrite acceptance of a norm like EN. I will argue that the instrumentalist cannot successfully carry out this task.

This argumentative strategy against the instrumentalist will perhaps appear odd given that we haven't yet established EN is a correct norm of belief. At this stage in the dialectic, the instrumentalist could happily concede that her view fails to accommodate a norm like EN; she will simply take this as reason for thinking EN is *not* an epistemic norm (or, alternatively, she will accept that it *is* an epistemic norm but insist that it only "applies" to a limited range of people). This is absolutely correct; my argument against instrumentalism will not be complete until I establish EN as a correct epistemic norm (a task that won't be carried out until Chapter 6). Thus, my full argument against the instrumentalist will not be complete until EN is established as a correct epistemic norm of belief. However, insofar as EN tracks our commonsense views about epistemic reasons for belief, demonstrating that instrumentalists are incapable of accommodating it will give us at least prima facie reason for rejecting the view. Once EN is established as a correct norm, however, the argument will be complete.

Let us move on, then, and consider different versions of the instrumentalist view. I would

like to consider four different varieties of instrumentalism.[96] I propose we sort different versions of the view according to two questions we can ask about the aims, interests, and goals invoked by the instrumentalist in her answer to the justificatory question vis-à-vis epistemic norms in general. As I've mentioned, the instrumentalist will argue that we should accept epistemic norms given their usefulness when it comes to the satisfaction of our aims, interests, and goals. Thus, according to the instrumentalist, for any individual S and epistemic norm N, the reason why S should accept N will ultimately have something to do with N's efficacy when it comes to the satisfaction of agential aims, interests, and goals.

However, as I've mentioned, there are two questions we can ask about these aims, interests, and goals: First, what is their *content*? *Cognitive* or *intellectual* goals somehow involve the attainment of knowledge or true belief. For instance, someone might have the goal of believing p if and only if p is true, or someone might want to know more about a certain historical event. Alternatively, there are *practical* goals (like getting to work on time). Thus, this is the first question we can ask about the relevant aims, interests, in goals; it concerns their content. Second, there is a question about *who has* the goal(s). As I mentioned, I will be concerned with *individualist* epistemic instrumentalism rather than *collectivist* epistemic instrumentalism. However, this doesn't mean that, in a case where S should accept some epistemic norm N, the reason for this *can't* involve the aims, interests, and goals had by someone

[96] My division of the instrumentalist position into these four groups is partly influenced by: Kornblith (2002, Chapter 5). I diverge with Kornblith only when it comes to the fourth variety of instrumentalism; a version which has only recently been articulated and defended. Also, Lockard (2013) is sensitive to a distinction between "intellectualist" and "pragmatist" forms of epistemic instrumentalism. However, Lockard does not distinguish between different versions of "intellectualist" epistemic instrumentalism. The versions of epistemic instrumentalism considered in sections 2-3 below count as "intellectualist" by Lockard's criteria, while the version considered in section 4 counts as "pragmatist".

other than S. Thus, we should be sensitive to a distinction between instrumentalist views which appeal to the aims, interests, and goals of the *individual agent* who has reason to accept some epistemic norm N (i.e. S herself), and instrumentalist views which appeal to the aims, interests, and goals of *other agents.* The four varieties of instrumentalism that I will consider are: Individual agent S + cognitive goal (sec. 2), other agents + cognitive goal (sec. 3), individual agent S + the fact that S has goals (sec. 4), and individual agent S + S's goal of φ-ing for reasons (sec. 5).

2. Individual Agent S + Cognitive Goal

The first version of instrumentalism appeals to a cognitive goal that individual agents possess. This version of instrumentalism is the subject of Thomas Kelly's (2003). According to this view, our acceptance of epistemic norms is justified given a *very general* cognitive goal pertaining to the acquisition of true beliefs and avoidance of false beliefs; one which is attributable to *all* agents. In other words, according to this view, whenever S has reason to accept some epistemic norm N, this is because of some very general cognitive goal that is, in fact, attributable to S. Thus, individuals will have reason to accept EN in virtue of their possession of this general cognitive goal. Kelly himself thinks that this view is mistaken, and his (2003) is an extended argument against it. Nevertheless, Kelly develops the view on behalf of the instrumentalist as a natural response to a certain worry. In presenting this worry, Kelly relies on intuitive observations regarding the "intersubjective" nature of epistemic reasons for belief such as the following:

> If both of us know that all of the many previously-observed emeralds have been green, then both of us have a strong reason to believe that the next emerald to be observed will be green, regardless of any differences which might exist in our respective goals. Similarly in arguing for my conclusions in this paper, I think of myself as attempting to provide strong reasons for believing my conclusions, and *not* as attempting to provide

> strong reasons for believing my conclusions for those who happen to possess the goal of the right sort.[97]

Kelly suggests that such observations about epistemic reasons are commonplace. This, however, seems to present a problem for the instrumentalist. While Kelly makes his point in terms of epistemic reasons, we could make a similar point regarding epistemic norms. Thus, it doesn't seem that epistemic norms fail to apply to some agent S simply because S doesn't care about a particular question or subject matter. If Kelly is right, then the application of epistemic norms (including EN) does not seem to depend on our adopting and retaining particular goals.

As Kelly points out, however, the instrumentalist has a possible response at this juncture. The instrumentalist could try to accommodate the data adduced by Kelly by attributing to *all* agents a very wide cognitive goal. By a "wide" cognitive goal I mean a goal to believe true propositions *in general*, not just true propositions *about matters that interest us*[98]. The instrumentalist could thus try to account for the data adduced by Kelly in the following manner: In the context of ordinary life we certainly *treat* epistemic reasons and epistemic norms as if they don't depend upon the aims, interests, and goals of agents, but this is only because the goal that *does* account for epistemic reasons and norms is universally shared. Compare: we tend to treat considerations related to the prolonging of human life in a certain way. When we present S with a reason to Φ where Φ-ing would prolong S's life (e.g. to stop drinking a certain liquid because, unbeknownst to S, it contains a lethal poison) we don't present this as a reason that S has *insofar as she has the goal of living a longer life*. However, one could argue that this is only because the

[97] Kelly (2003, p. 621).

[98] The instrumentalist view currently under consideration bears some resemblance to Friedman's "zetetic" approach to epistemic norms, but there's an important difference: The view currently under consideration involves attributing to individuals a very *general* cognitive goal related to the attainment of true beliefs or accurate information, i.e. a goal that isn't confined to certain specific topics or subject matters (e.g. how many windows the Chrysler Building has).

goal of living longer is so close to universally shared that we simply take it for granted that any particular individual possesses it. An instrumentalist could try to make a similar move in response to Kelly's observations. Specifically, she could try to accommodate the observations about the apparent "intersubjective" or "non-hypothetical" nature of epistemic reasons and norms by appealing to some aim or goal that *isn't* idiosyncratic. Thus, the goal would have to be universally shared and it would have to pertain to the acquisition of truth and the avoidance of falsity *in general.* This, then, is a move that a proponent of "individual agent S + cognitive goal" could make in response to Kelly's considerations

While this is certainly a maneuver the instrumentalist can make at this juncture, as Kelly points out, people in general do not possess a goal of this kind.[99] While it is perhaps plausible to suggest that we possess relatively wide cognitive goals, there are very real limits to how wide even our widest cognitive goals are. For instance, who among us has a cognitive goal that would be satisfied by learning that the 78th person in the Bloomington, Indiana telephone book has a last name that contains an odd number of letters? Examples of such "trivial" truths abound. Nevertheless, we *can* come to have conclusive evidence for these trivial truths. Moreover, in at least some cases where such evidence is brought to our attention, it is not implausible to suggest that we thereby have a reason to believe the relevant proposition (as is implied by EN). For instance, Kelly notes that

[99] This is, of course, an empirical question at the end of the day, but the claim that *all* individuals possess such a goal is clearly doubtful. Hilary Kornblith, an avowed instrumentalist, puts the point nicely: "It is important to note…that any attempt to gain universal applicability by appeal to goals that all humans in fact have will almost certainly run afoul of the facts. Human beings are a very diverse lot; some of us are quite strange. It is hard to imagine making a plausible case for any particular goal or activity that is genuinely universally valued." Kornblith (2002, p. 150).

> from the fact that some subjects are matters of complete indifference to me, it does not follow that I will inevitably lack epistemic reasons for holding beliefs about those subjects.[100]

Thus, Kelly concludes, the wide cognitive goal maneuver fails to save the instrumentalist. People do not possess the kind of goal that is needed in order to do the requisite work.

Kelly's points here strike me as correct so I won't spend any more time on his response to this particular instrumentalist maneuver[101]. What I'd like to do instead is take stock and look forward. As we have seen, Kelly argues that the wide cognitive goal maneuver fails to save "individual agent S + cognitive goal"; people in general do not possess the kind of goal that is needed in order to accommodate Kelly's observations. It seems, then, that there will be instances in which EN applies to a person even though that person *doesn't* possess a cognitive goal of the right kind. For instance, we might still expect someone to treat some strong evidence E for the truth of p as a warranting reason to believe p *even if* they lack some specific cognitive goal that would be satisfied by believing p. If this observation is on the right track, and if "individual agent S + cognitive goal" is the only available form of instrumentalism, then it looks like instrumentalism faces insuperable challenges. Specifically, it looks like the only line of response available to the instrumentalist is the maneuver described above, viz. the attribution of a very wide cognitive goal to all subjects. But, as we have just seen, this is clearly implausible.

Fortunately for the instrumentalist, "individual S + cognitive goal" is not the only version of the view. I will consider more instrumentalist views below. Also, at this stage of the dialectic,

---

[100] Kelly (2003, p. 625).

[101] Note that, in saying that Kelly's comments strike me as correct, I am of course not avowing commitment to the view that I labeled as "evidential minimalism" in chapter 1. What strikes me as correct is the following: Epistemic norms don't to fail to apply to us simply because we lack certain cognitive goals. There are different ways of trying to accommodate this thought. The minimalist position is one possible route, in the next chapter I will begin to develop an alternative view.

one might suggest abandoning the hope of establishing a wide scope of application for a norm like EN; perhaps when a person doesn't possess the requisite goal(s) the norm *doesn't* apply her, contrary to appearances. As I have mentioned, my argument against the instrumentalist view won't be complete until I establish EN as a correct epistemic norm, so this is an avenue that is available to the instrumentalist at this point. However, insofar as EN tracks our ordinary commitments regarding epistemic reasons for belief, the above gives us at least prima facie grounds for rejecting "individual S + cognitive goal".

3. Other Agents + Cognitive Goal

The second variety of instrumentalism also appeals to cognitive goals. However, unlike the first variety, it does not rely on an implausible claim regarding the attribution of a very wide cognitive goal to all agents. Instead, this view appeals to cognitive goals that are had by only certain agents. According to this version of instrumentalism, the reason why some agent S should accept an epistemic norm N can involve the goals had by somebody *other than S herself*. There are two ways of developing this version of instrumentalism. On the first way, the other agents appealed to are *actual*, and on the second way they are *counterfactual*. Let's start with the former.

According to this version of the view, the reason why S should accept some epistemic norm N will involve the aims, interests, and goals of those individuals who *do*, as a matter of fact, take on a cognitive goal of the relevant sort, viz. a goal of attaining truth and avoiding falsity[102]. This version of instrumentalism thus goes some way towards accommodating EN

[102] I am not aware of anyone who has defended this version of epistemic instrumentalism. However, Kate Manne has defended a view about practical normativity that bears some resemblance to this approach (Manne (2016)). According to Manne's "democratized" Humean position, an agent's reasons for action may be explained by the desires had by individuals other than that agent herself.

without also falling prey to the problem confronting "individual agent S + cognitive goal". After all, many people *do*, in fact, value certain intellectual pursuits such as scientific inquiry (whether as practitioners or as non-practitioners). There is a question about whether or not the aims or goals of, say, scientific inquiry are "wide" enough in order to accommodate the worry sketched in the last section. However, let's set that question to the side and assume for the moment that the goal is wide enough to assuage that worry.[103] Thus, according to this version of "other agents + cognitive goal", whenever an agent S has reason to accept some epistemic norm N, this is because of the cognitive goal(s) had by those individuals who value certain intellectual pursuits (such as scientific inquiry) whose characteristic aim is the attainment of truth and avoidance of falsity.

Notice that, in any given case where S has reason to accept some epistemic norm N, S *herself* may or may not possess the relevant cognitive goal. An issue for this version of instrumentalism arises when we consider cases where S herself does not possess the relevant cognitive goal. So imagine some agent S who does not value scientific inquiry or any other intellectual pursuit whose aim is the attainment of truth. When it comes to her doxastic life S is primarily concerned with taking on beliefs that corroborate her pre-existing biases and promote her own selfish interests. As it turns out, in order for S to achieve these latter aims, she sometimes has to take on and/or maintain beliefs that are at odds with the evidence that she possesses. Imagine a scenario in which S possesses strong evidence E which tells in favor of the truth of p but believing in accordance with this evidence wouldn't satisfy any goal that S has. Also say that, at time t, the evidential connection between E and p is clear to S. According to EN, S has a normative (warranting) reason to believe p at t. According to "other agents + cognitive

[103] I will return to this issue momentarily.

goal" S should accept this norm because there are some *other* individuals who have a goal that would be furthered by S's doing so. The thought is that, by accepting a norm like this, S will be increasing the likelihood that she will acquire true beliefs and avoid false beliefs, and this will help further the cognitive goal(s) of other people.

There are several problems with this account. I would like to focus on just two. First, let's assume that S's acceptance of the norm actually *would* further the goals that are actually taken on by other individuals. Even assuming this, there's still the following issue: Why does this generate a reason *for S* to accept the norm? Say there's a group of people who share a goal related to the acquisition of truth and avoidance of falsity. Let's even say that this is a fine goal to have; the acquisition of truth and the avoidance of falsity is a worthwhile pursuit. However, it doesn't automatically follow from this that people who *aren't* in this group have reason to facilitate satisfaction of the group's aims or goals. Groups of individuals can take on various worthwhile pursuits. A group of archeologists might convene with the shared aim of uncovering a certain ancient artefact, a group of astronomers might work together in order to study some distant star or planet, a group of scientists might collaborate with the aim of figuring out whether or not a new drug is effective in treating a disease, etc. Again, these all seem to be worthwhile pursuits. However, it doesn't follow from this that individuals who *aren't* in the various groups have reason to facilitate satisfaction of their aims. Thus, even if we grant that S's acceptance of some epistemic norm N would further the goals taken on by a group of individuals that S doesn't belong to – a group of individuals who have the shared aim of acquiring truth and avoiding falsity – it's not clear that this would generate a reason *for S* to accept the norm.

However, even if this issue could somehow be addressed, there's still another problem confronting "other agents + cognitive goal". Let's now *grant* that S has reason to facilitate the

satisfaction of the cognitive goal(s) taken on by the members of certain groups that S doesn't belong to. Even if there are other individuals who have certain cognitive aims or goals (i.e. goals related to the acquisition of truth and avoidance of falsity), these aims and goals surely do not amount to *unrestricted* goals related to the acquisition of *any truth whatsoever*. Rather, as we saw above, the aims or goals of groups interested in the truth are focused and circumscribed in various ways. For instance, research groups tend to be focused on a particular question or subject matter. Return again to S; the person who doesn't value scientific inquiry and who merely seeks to take on beliefs that promote her own interests and corroborate her pre-existing biases. Imagine that S possesses (and can clearly see the evidential import of) strong evidence E for the truth of some trivial or insignificant claim p, e.g. that there are 254 specks of dust on my computer screen right now. According to EN, S has a normative (warranting) reason to believe this claim. According to "other agents + cognitive goal", S has reason to accept this norm since doing so would promote the satisfaction of the cognitive goal(s) had by a certain group of individuals, viz. those individuals who value intellectual pursuits whose characteristic aim is the attainment of truth. However, given that the aim of such pursuits is to uncover only a *circumscribed* set of truths (i.e. those that are somehow *significant* or *of interest*), it does not look as though the proponent of "other agents + cognitive goal" can secure S's acceptance of EN.

Recall, however, that there's another way of developing "other agents + cognitive goal". The way we've just been considering appeals to other agents that are *actual*. There is another way of developing this view that appeals to other agents that are *counterfactual*. So imagine a version of "other agents + cognitive goal" that appeals to *idealized* agents. The idea here is that, for any actual agent S, there is an idealized version of S who can somehow set normative standards for her actual self. What makes an agent "idealized" will vary depending on the

particular view, but such views usually require idealized agents to be operating with accurate information and an undistorted motivational structure.[104] Such an approach could still be in line with the general thrust of the instrumentalist view. For instance, we could say that the *starting set* of an individual's aims, interests, and goals will be her *actual* aims, interests, and goals. The idealization process could then be constrained by the individual's actual psychology; it will take as input the agent's actual aims, interests, and goals. Thus, according to the idealized agent version of "other agents + cognitive goal", whenever there's some epistemic N that S should accept, this is because accepting N would satisfy or promote some cognitive goal that S *would* have, were she suitably idealized.

The problem for this view is that, even when agents are idealized, they will surely come to have various *different* cognitive goals. Perhaps, once suitably idealized, every agent would come to have *some* cognitive goal or other. But these goals will no doubt vary, and many of them will surely fall short of an *unrestricted* goal to believe the truth about any subject matter whatsoever. Given this, the idealized agents version of "other agents + cognitive goal" will also fail to secure a sufficiently wide scope of application for a norm like EN; even suitably idealized agents will often fail to have the kind of cognitive goals that would be needed in order to underwrite acceptance of a norm like EN.

Note that, once again, a proponent of "other agents + cognitive goal" could acknowledge these issues but nevertheless persist in defending the view, in either of the forms presented here. For instance, one could deny that EN actually *is* a correct epistemic norm, or one could maintain

[104] What makes an agent's motivational structure "undistorted" will also vary depending on the particular view. The *locus classicus* for views in this neighborhood is perhaps Brandt (1979). According to Brandt, some consideration constitutes a reason for some subject to Φ only if Φ-ing would promote a goal or desire that this subject *would* have after undergoing cognitive psychotherapy. Thus, on Brandt's view, an idealized agent's motivational structure counts as "undistorted" in virtue of having survived cognitive psychotherapy.

that it is but hold that it only "applies" to a limited range of people. At this stage in the dialectic, these options are still available. However, insofar as EN tracks our ordinary commitments regarding epistemic reasons for belief, the above gives us prima facie reason for rejecting the view. Let us move on, then, and consider other instrumentalist possibilities.

4. Individual Agent S + The Fact that S has Goals

So far we've only looked at varieties of instrumentalism that appeal to "cognitive" or "intellectual" goals. Hilary Kornblith has defended a version of epistemic instrumentalism which, similar to our first variety, appeals to the goals of the individual agent S who has reason to accept some epistemic norm N. However, instead of appealing to some *particular* goal had by S, Kornblith's instrumentalism appeals to the fact that S has any goals whatsoever, whether cognitive or otherwise.[105] Kornblith's core insight is that, *whatever* an individual's aims, interests, or goals happen to be, it will be beneficial for that individual, on the whole, to be responsive to the evidence. This is because doing so will conduce to the attainment of that individual's goals, whatever they might be.

Consider the following example offered by Kornblith: Imagine we're trying to buy a toaster. In choosing between two different toasters, we have to engage in a cost-benefit analysis; i.e. we have to figure out the consequences of each potential purchase, assign values to each of these outcomes, and perform some calculations in order to determine which course of action will best serve our interests. If we possess cognitive systems that give us true beliefs, then, through performing this analysis, we will come to find out which decision will, in fact, best serve our interests. This will then allow us to actually satisfy our aims in this case, and similarly in other cases where we're trying to achieve our desired ends. Alternatively, if we possessed cognitive

[105] Kornblith (2002, Chapter 5).

systems that were, say, *happiness-conducive*, then we'd be hampered in our efforts to achieve our own ends. As Kornblith puts is,

> Allowing our cognitive systems to be determined by the totality of our interests exclusive of the truth thus undermines our ability to make choices, outside the cognitive realm, that are conducive to those very interests…It seems that someone who cares about acting in a way that furthers the things he cares about, and that includes all of us, has pragmatic reasons to a favor a cognitive system that is effective in generating truths, whether he otherwise cares about the truth or not.[106]

Kornblith argues that, on the basis of these considerations, we are instrumentally justified in accepting epistemic norms. Thus, on the Kornblithian approach, whenever a subject S has reason to accept some epistemic norm N, this will be because doing so will help promote S's aims, interests, and goals *on the whole and in the long run*. On Kornblith's approach, then, epistemic norms could be thought of as rules that we're instrumentally justified in accepting. Even though, on some particular occasion, following a certain rule might not satisfy a specific goal actually had by S, S could still be justified in accepting the rule given that doing so is a good way to satisfy her aims, interests, and goals on the whole and in the long run.

In order for the present version of instrumentalism to secure a wide scope of application for a norm like EN, it has to be the case that a *very large class* of individuals has aims, interests, and goals that would be best served, on the whole, by accepting a norm that *always* treats strong evidence as a normative (warranting) reason for belief when that evidence is possessed, and its evidential import is clear. But this is highly doubtful. Imagine again the agent that I described in the previous section who does not value scientific inquiry and only seeks to take on beliefs that corroborate her pre-existing biases and further her own selfish interests. Let's also say that this agent's goals are fairly fixed in the sense that they will not alter radically over time. Would this

[106] Ibid. p. 155.

agent's aims, interests, and desires be best served, on the whole, by adopting and following a norm like EN? Certainly many of her aims, interests, and goals will only be satisfied if she displays *some* sensitivity to her evidence; I think Kornblith's points above suffice to establish this. But it also seems overwhelmingly likely that there will be many cases where this agent's aims, interests, and goals will not only not be *furthered* by believing in accordance with her evidence, but will actually be *thwarted*. It thus seems that, for such an agent, a norm like EN will not best serve her overall interests. It seems, instead, that this agent's overall interests would be better served by adopting and following a more restricted evidential norm, e.g. one which only treats as a warranting reason for belief evidence that somehow bears on her pre-existing (narrow-minded) aims, interests, and goals.

One might push back here by noting that "warranting reasons" are not very demanding. Recall that warranting reasons are considerations that it would be appropriate to base one's beliefs upon, whether or not one actually does so believe. Given this, it may seem less plausible that the above agent's aims, interests, and goals wouldn't be best served by adopting and following a norm like EN. After all, one could "adopt and follow" such a norm by cherry-picking the beliefs that one actually takes on so as to acquire the goal-advancing ones while avoiding the goal-thwarting ones. However, if the above agent is anything like the rest of us when it comes to her belief-forming mechanisms, then she does not possess the ability to cherry-pick in this way. Moreover, even if she *did* possess this ability, this wouldn't give us any reason to think that her aims, interests, and goals would be best served by accepting EN. Rather, given the nature of her aims, interests, and goals, it seems that a more restricted doxastic norm would better serve her ends.

One could grant the points just made while nevertheless insisting that the agent that I

describe above is merely possible and not actual. In other words, according to this response, there are no agents with the narrow-minded aims, interests, and goals described above, or with aims, interests, and goals that are sufficiently similar. In response, I would first like to reiterate that the instrumentalist here shoulders an enormous dialectical burden. As Kornblith himself stresses, human beings are "a very diverse lot" and "some of us are quite strange".[107] It seems rather implausible that there are *no* agents with aims, interests, and goals like those had by the agent described above. This is, however, an empirical question at the end of the day. Nevertheless, notice that, while many of us are (hopefully) not as narrow-minded as the agent described above, we do sometimes acquire evidentially supported, true beliefs that thwart some of our goals and interests (imagine, once again, the individual in MOVIE SPOILER). Additionally, there is empirical evidence which suggests that human beings are psychologically constituted such that operating in accordance with the unrestricted EN would be detrimental when it comes to the overall achievement of our personal aims, interests, and goals. For instance, Taylor and Brown (1988) note that human subjects frequently manifest unrealistically positive views of the self, exaggerated perceptions of personal control, and unrealistic optimism. Not only have such tendencies been found to be ubiquitous in human subjects, but they also seem to play an important role when it comes to overall functioning:

> …a great deal of research in social, personality, clinical, and developmental psychology documents that normal individuals possess unrealistically positive views of themselves, an exaggerated belief in their ability to control their environment, and a view of the future that maintains that their future will be far better than the average person's. Furthermore, individuals who are moderately depressed or low in self-esteem consistently display an absence of such enhancing illusions. Together, these findings appear inconsistent with the notion that accurate self-knowledge is the hallmark of mental health.[108]

It seems likely, then, that there are certain standing beliefs that we (perhaps unconsciously) take

[107] Kornblith (2002) p. 150.
[108] Taylor and Brown (1988) p. 197.

towards ourselves which are not well supported by the evidence yet also play an important role when it comes to the achievement of our aims. We can imagine an agent who, upon encountering evidence for the falsity of one such belief, revises her beliefs about herself and thereby negatively affects her ability to achieve her aims. Say, for example, this agent comes to find out that she's not quite as good as she thought she was at chess, which is a cherished pastime of hers. This, in turn, makes her anxious during future gameplay and hinders her ability to enjoy chess.

I take these examples to provide us with good grounds for thinking that at least *some* of us are such that our aims, interests, and goals would not be best served, on the whole, by accepting a norm like EN. If this is correct, then the instrumentalist has, once again, failed to offer a view which establishes a sufficiently wide scope of application for EN. We thus find ourselves in the same dialectical situation once more. We can either maintain the instrumentalist view on offer (thereby rejecting EN as an epistemic norm), or we can try to develop the view again in the hope of securing a wide of application for EN. Let us consider one more instrumentalist attempt to carry out this task.

5. Individual Agent S + S's Goal of Φ-ing for Reasons

The last version of instrumentalism that I would like to consider has been articulated and defended by Côté-Bouchard (2015). Côté-Bouchard relies on the observation that we all, as agents, act, feel, and believe for at least *apparent* reasons. By an "apparent" reason he means a consideration that we *take* to be a "genuine" reason. "Genuine" reasons, according to Côté-Bouchard, are facts or true propositions. Thus, a reason that p is a genuine (and not merely apparent) reason only if p is true. If p is false and some subject Φs with p as her motivating reason, then she has Φ-ed for a merely apparent reason. Côté-Bouchard argues that whenever we Φ for an at least apparent reason that p, we *believe* that p. According to Côté-Bouchard:

> it would be absurd to assert, e.g. , 'my reason for bringing my umbrella is that it is raining outside, but it is not raining outside'. It would be equally absurd to assert, e.g., 'Myriam's reason for being angry is that someone stole her cupcakes, but she does not believe that someone stole her cupcakes.[109]

Côté-Bouchard's suggestion is that, whenever we Φ for an at least apparent reason, we aim at Φ-ing for a genuine reason; a goal that can only be satisfied by acquiring true beliefs. The best way to satisfy this (universally shared) goal, moreover, is to accept doxastic norms that instruct us to believe on the basis of our evidence. Thus, according to Côté-Bouchard, whenever S has reason to accept some epistemic norm N, this will be because S has the goal of Φ-ing for a genuine reason whenever she Φs for an at least apparent reason, and acceptance of N represents an effective way to promote or satisfy this goal.

Notice, first, that there is a way to satisfy the goal proposed by Côté-Bouchard that does not involve accepting doxastic norms that instruct us to believe in accordance with our evidence. The goal is formulated, once again, as follows: Φ for a genuine reason whenever you Φ for an apparent reason. One way to satisfy this goal would be to *never* Φ at all, even for apparent reasons. Compare: I have the goal of going to the doctor whenever I get sick. One way to satisfy this goal is to take measures to never get sick. But let's set this worry to the side. The real issue for the present variety of instrumentalism is that, like Kornblith's view, it cannot secure a wide scope of application for a norm like EN. Let's grant that all agents have the goal of Φ-ing for genuine reasons whenever they Φ for merely apparent reasons. Let's also grant that, in order to satisfy this goal, agents have to display some sensitivity to their evidence when forming and maintaining doxastic attitudes. In order for the present variety of instrumentalism to secure a wide scope of application for a norm like EN, it has to be the case that a *very large class* of

[109] Côté-Bouchard (2015, pp. 350-51).

individuals are such that their *individual* goals of Φ-ing for genuine reasons whenever they Φ for merely apparent reasons would be best promoted by accepting a norm like EN. This, however, is highly unlikely. Imagine someone who only needs evidentially supported beliefs about a limited range of topics in order to satisfy her goal of φ-ing for genuine reasons whenever she φ's for merely apparent reasons. Perhaps this person's life is fairly boring and predictable; she exercises her agency by acting and thinking, but only in fairly routine ways. When this person comes to possess strong evidence that has no bearing whatsoever on her pre-existing aims and interests – say, evidence that it rained in Paris on October 5th, 1643 – will that evidence constitute a warranting reason for her? EN says that it will, but the present variety of instrumentalism is unable to accommodate this. This person has the goal of φ-ing for genuine reasons whenever she φ's for apparent reason, but for this particular agent that goal would be better satisfied by adopting and following a more restricted doxastic norm.

We have, once again, failed to secure a wide scope of application for EN on instrumentalist grounds. We are thus in the same dialectical situation once more: We can accept the current instrumentalist view on offer (thereby rejecting EN as an epistemic norm), or we can try to develop the instrumentalist view once more in an attempt to accommodate EN. At this point I will leave it open whether there are further ways to develop the instrumentalist position. My aim was to survey some of the main varieties of the view, but there are perhaps further ways of developing the position that I have not considered here. I would like to conclude instead with a few comments about the dialectic as it currently stands.

6. Conclusion

My aim in the present chapter was to consider different ways of understanding epistemic norms as instrumental norms. I have argued against four different varieties of the instrumentalist

view. My motivation in considering these views stems from considerations presented in the last chapter. As I noted there, my overarching goal is to understand evidential normativity. I argued that the best way to do so is by considering which evidential norms are correct epistemic norms of belief. In particular, I organized much of the discussion around the evidential norm EN. My proposal was that we begin with the justificatory question about epistemic norms for belief in general. However, in order to ask this question, we need some specification of the "epistemic" itself. In the previous chapter I considered and rejected Friedman's account which would have us collapse the epistemic with the zetetic. However, one might feel that something about Friedman's approach is on the right track. Specifically, one might attempt to appeal to a wider set of aims, interests, and goals, not just those unique to the activity of inquiry, in an attempt to understand epistemic norms. I considered four different attempts to pursue this idea in this chapter, and I've argued that they all face significant challenges. A recurring issue was that these views do not seem to possess the resources to accommodate a wide scope of application for a norm like EN. Insofar as EN tracks our ordinary commitments regarding epistemic reasons for belief, this gives us prima facie reason for rejecting epistemic instrumentalism. After EN is established as a correct epistemic norm, the argument against epistemic instrumentalism will be complete.

If Friedman's zetetic approach and epistemic instrumentalism both face significant challenges, then we're still left with the following question: How are we to understand "epistemic" norms of belief? In the next chapter I will begin to develop my preferred approach to epistemic norms. As we will see, my approach will locate epistemic norms of belief within our social practices of interpersonal accountability. I will argue that, by pursuing matters in this way, we will be able to develop a satisfactory answer to the justificatory question vis-à-vis epistemic

norms in general, and evidential norms in particular. Once this answer has been articulated and defended, we will be able to develop an answer to the content question vis-à-vis epistemic norms of belief.

Chapter 4: Epistemic Accountability

1. Introduction

In the previous chapter I argued against various instrumentalist attempts to understand epistemic norms. According to epistemic instrumentalists, we should accept epistemic norms since doing so represents an effective way to satisfy our aims, interests, and goals. Each variety of the instrumentalist view faced significant challenges. Specifically, each view failed to secure a wide scope of application for a norm like EN. However, I also noted that an instrumentalist might accept these results and simply *deny* that EN is an epistemic norm of belief. Given the "Application-Acceptance Connection" (AAC), this means that the instrumentalist must hold that, in cases where S's acceptance of EN cannot be secured on instrumentalist grounds, the norm does not "apply" to S; i.e. it isn't appropriate for other people to *hold S accountable* for complying or failing to comply with the norm. In the present chapter, I will begin to explain why this position is untenable. Specifically, I will argue that, when we attend to our actual practices of interpersonal epistemic evaluation, we will be able to discern a practice of *epistemic accountability*. In subsequent chapters, I will argue that our practice of epistemic accountability is a *legitimate* social practice, and that it is structured by certain norms of belief (including EN). This will be the way that I intend to answer the justificatory question vis-à-vis epistemic norms in general (and evidential norms in particular): We should accept such norms given that they structure a legitimate social practice that we participate in. However, before carrying out these further tasks, we must first pursue a certain descriptive/interpretive project vis-à-vis our existing social epistemic practices. Specifically, we must establish that *there is* a distinctly epistemic form of interpersonal accountability. This is the task that I will carry out in the present chapter.

Recall that my overarching aim is to understand evidential normativity. I have argued that the best way to pursue this question is by considering which evidential norms are correct epistemic norms of belief. I also proposed beginning with the justificatory question vis-à-vis epistemic norms in general. In chapter 2, I noted a challenge that confronts the approach that I've adopted here related to the very idea of the "epistemic". In order to ask the justificatory question about epistemic norms in general, we need some specification of the "epistemic" itself. However, there doesn't appear to be a non-arbitrary way to demarcate the class of "epistemic" norms from other, non-epistemic, norms.

My view is that understanding our actual practices of interpersonal epistemic criticism and evaluation represents the best way to overcome this difficulty. The basic idea is as follows: Norm violations can be met with certain kinds of *responses and reactions* on the part of other people. For instance, if you break a promise to me without a good excuse, it seems that I can appropriately *hold you responsible* in certain ways, e.g. by blaming you or by modifying my level of trust in you. Moreover, it seems that certain classes of norms have *distinctive kinds* of reactions that are called for when they are violated, e.g. blame in the case of moral norms, or sanctions in the case of legal norms. My thought is that, if we can locate a distinctive kind of response that can be appropriately taken up by others in reaction to a person's doxastic transgression, then that will help demarcate the class of epistemic norms.[110] As I mentioned

[110] Doesn't this presuppose that we *already know* which norms are the "epistemic" norms? In other words, doesn't this presuppose that we've already answered what I've called the "content" question vis-à-vis epistemic norms? Here I'll echo a point I made in chapter 2: I take it as obvious that, in the context of ordinary life, we hold people to certain expectations when it comes to the things that they believe, e.g. we admonish things like wishful thinking, believing on the basis of insufficient evidence, and believing bald contradictions. We can understand these norms or standards as the "starting materials" for our investigation. This doesn't stack the deck

above, approaching matters in this way could then help us answer the justificatory question about epistemic norms in general, and evidential norms in particular.

The present chapter is dedicated to establishing the claim that we do, in fact, have a social practice of epistemic accountability. As I have mentioned, before we can turn to questions about the legitimacy of such a practice, we must first establish that such a thing even exists. In the next section (sec. 2). I will begin carrying out this task. In sec. 2 I will introduce some conceptual distinctions from the literature on moral responsibility in order to help frame the subsequent discussion. In sec. 3, I will turn to the question of epistemic accountability. I will develop three criteria that a form of response R must meet in order to count as an epistemic accountability response. I will argue that there is a form of response at play in our actual practices that meets these criteria. Sec. 4 is devoted to establishing that this form of response meets the last of the three criteria, viz. that the response goes beyond mere assessment vis-à-vis a standard. I conclude in sec. 5.

2. Accountability, Attributability, and Answerability

To help structure the discussion, I will be relying on some conceptual distinctions from the literature on moral responsibility. While these conceptual distinctions are usually invoked in discussions concerning practical agency, I think their incorporation into the epistemic domain will prove to be fruitful. However, we must also proceed with caution. Specifically, we should be careful not to assume at the outset that our practice of epistemic accountability (if such a thing

in favor of certain philosophical views over others: these materials can be developed in various theoretical directions, and they are amenable to different conceptions of the "epistemic".

exists) mirrors our accountability practices in other domains[111]. Keeping this in mind, let us proceed.

The notion of "accountability" is usually situated within a broader cluster of concepts relating to agential responsibility. Talk of "responsibility" in ordinary life is complicated by the fact that we sometimes attribute responsibility to inanimate objects and the mechanical proceedings of nature, e.g. "the avalanche is responsible for the destruction of the cabin". What we usually have in mind here is *causal* responsibility. However, when it comes to *agential* responsibility, we usually have in mind something distinct from mere causal responsibility. Imagine, for instance, a person who has a seizure in public and knocks over and breaks an expensive vase. While there might be a causal story to tell about the breaking of the vase, the person in question is, in an important sense, not responsible for its breaking. Consider, by way of comparison, someone who picks up the vase and intentionally smashes it on the ground out of malice. This latter person clearly *is* responsible for its breaking. The attribution of responsibility in this latter case is tied to our blaming practices; it would be appropriate in the latter case, but not the former, for other people to *blame* the person for their conduct. This might involve various attitudinal responses (e.g. anger, resentment, indignation), expressions of these attitudes (perhaps directly to the offender in the form of protest), requests for recompense, demands for apology, etc. Responding to an individual in these latter ways amounts to "holding" a person responsible; such responses involve more than mere negative appraisal or assessment vis-à-vis a standard.

[111] To clarify: At this point, I do not mean to rule out the possibility that epistemic accountability simply reduces to, say, moral accountability. For instance, we may find that the way in which we hold each other accountable for the things that we believe doesn't support a *distinctly epistemic* form of accountability. My point is that we shouldn't *assume* this at the outset.

Imagine, for instance, a "cool" or stoic observer of the vase smasher; someone who recognizes that a bad act has been done and that it issued from an objectionable character, but who doesn't take up any of the further responses listed above. For example, they aren't emotionally exercised by the act, they express no anger or protest, they demand no apology, etc. Such a person might be engaged in a ("detached") form of interpersonal assessment or evaluation, but they wouldn't be *holding the person responsible* for their act[112].

This latter dimension (or "face") of agential responsibility was identified by Gary Watson as responsibility as *accountability*.[113] Responsibility as accountability thus concerns the conditions under which it's appropriate to subject an individual to certain *reactions* in response to her thought and conduct, e.g. blame (which may include reactive attitudes such as resentment) and sanctions of various kinds. When describing the forms of accountability that were appropriate to take up in response to the vase smasher, I had in mind *moral* accountability, especially *blame*. It's important to note, however, that there are other forms of accountability. For instance, there are *legal* forms of accountability, as well as *institutional* and *professional* forms of accountability (among others). Consider, for instance, the following example from Kauppinen (2018):

> DRUNK DRIVER: [S]uppose it is necessary for complying with the norm of avoiding reckless driving that your blood alcohol level is below 0.1% by volume. You've had a few drinks that a reliable friend told you had very low alcohol content, and you don't feel

[112] Note that this point holds even for views which downplay the role of the reactive emotions when it comes to moral blame (e.g. the view of Scanlon (2008)). On Scanlon's view, moral blame involves *modifying one's relationship* vis-à-vis another. The "cool" or stoic observer described above might engage in "detached" appraisal or evaluation without modifying or reconfiguring his relationship with another person.

[113] Watson (2004). For a recent discussion regarding Watson's distinction between responsibility as "accountability" and responsibility as "attributability", see: Wolf (2019). I will discuss the notion of attributability momentarily.

> drunk, so that you're highly justified in believing that your blood alcohol level is below the required level, and you decide to drive home. But in fact you are over the limit, drive erratically, and are caught by the police. In this case, you may well have an excuse from moral sanctions — in the circumstances, it wouldn't be fair to resent you, say. But you plausibly don't have an excuse from legal sanctions — you should pay your fine. The legal sanction doesn't imply that you're morally bad, or that there's something wrong with your character or your will.[114]

As Kauppinen explains, it doesn't seem fair to subject the individual in this case to certain *moral* forms of accountability (e.g. blame responses including resentment). However, it *does* seem fair to subject them to certain *legal* forms of accountability (e.g. fines). In the case of DRUNK DRIVER, the inappropriateness of moral accountability seems to be related to the fact that the individual's conduct doesn't disclose or reveal anything morally objectionable about their commitments, evaluative orientation, or character; they were simply non-culpably ignorant of something (the strength of their drinks) which led to a certain dangerous act (erratic driving). This is in stark contrast to the case involving the malicious vase smasher where the action *did* reveal something morally objectionable about the person's values and commitments. The case of DRUNK DRIVER shows us that, even if *moral* accountability is inappropriate in a certain case, there still might be other, non-moral, forms of accountability (e.g. legal) which *are* appropriate. Moreover, even in instances where someone's conduct *does* disclose or reveal something morally objectionable about their commitments, evaluative orientation, or character, holding them morally accountable via blame responses might *still* be inappropriate. Consider, for example, instances where a person loses "standing" to blame[115]. As Lacey and Pickard (2021)

[114] Kauppinen (2018) pp. 8-9.

[115] This isn't the only kind of case where a person acts in a way that discloses something morally objectionable about their commitments, evaluative orientation, and character, yet moral accountability is either altogether inappropriate or at least mitigated. Other possible cases involve individuals who acquire morally objectionable character traits through no fault of their own (e.g. by being raised in a certain way). Such traits might dispose a person to engage in

note:

> Certain facts about the blamer may undermine their right to make the [moral] accusation, allowing for the accusation to be deflected without it being denied. Such challenges include the charge of hypocrisy ("look who's talking!"), complicity ("you're just as much to blame for this as I am!"), and meddling ("it's none of your business!"). These challenges attempt to rebuff the accusation by levying an accusation in turn, namely, that, in virtue of these facts about *them*, the blamer does something *wrong* in blaming, which undermines their right to censure the wrongdoing and demand the wrongdoer answer to the charge.[116]

Return again to our vase smasher. Say that I'm also a malicious vase smasher. In virtue of the fact that I'm guilty of the same sort of misconduct, it seems that I have lost the standing to engage in certain forms of moral censure and criticism vis-à-vis our initial vase smasher; I would be guilty of hypocrisy were I to do so ("Who am I to judge?"). Even so, according to Lacey and Pickard, there are cases like this where someone loses standing to blame but nevertheless *retains* the authority to engage in certain non-moral forms of accountability[117].

Thus, our "accountability" practices (whether moral or non-moral) are clearly varied and complex. However, as I mentioned earlier, the rough idea is that responsibility as accountability concerns the conditions under which its appropriate to subject a person to certain kinds of responses and reactions, including (but not limited to) blame responses. There is another dimension (or "face") of responsibility that will also be relevant for our discussion regarding

---

immoral conduct, but, given the etiologies of their characters, certain accountability responses might be inappropriate or at least mitigated.

[116] Lacey and Pickard (2021) p. 267.

[117] Consider Lacey and Pickard's memorable example: Imagine that a parent goes out and leaves their older son in charge of their younger daughter with clear instructions to send the younger daughter to her room if she misbehaves. Suppose that the older son then "winds up the younger as only older children know how to do." The younger then throws a tantrum and begins to wreck the house. Even though the older child is complicit in the younger's misbehavior, he still seems to retain the authority to hold her accountable by sending her to her room.

epistemic accountability. Watson (2004) distinguishes responsibility as *accountability* from responsibility as *attributability*. As we have seen, "accountability" involves various forms of interpersonal criticism, blame, and sanction. As a result, the appropriateness of these forms of response will be sensitive to considerations of *fairness.* We have already seen this in the case of moral blame; sometimes a person can lack "standing" to blame another as a result of hypocrisy, complicity, or meddling. In addition to this, certain forms of accountability might be inappropriate (or at least mitigated) in cases where a person lacks a certain amount of *control* over what they do, or where a person couldn't realistically *avoid* doing a certain thing.

Responsibility as *attributability* is a less stringent (yet still significant) form of responsibility; it concerns whether or not a person's thinking or conduct *expresses* or *discloses* their "true" or "deep" self. There are very difficult questions here about what it is to have a "self" in the first place, and what it takes to express or disclose one's ("true" or "deep") self. I won't be able to settle these questions here. It will suffice for my purposes if we can get the intuitive idea on the table. According to Watson, what's at issue when it comes to responsibility as attributability is "an individuals' fundamental evaluative orientation."[118] Watson describes responsibility as attributability as the "aretaic" face of responsibility since it often involves an individual's virtues and vices as manifested in thought and conduct. He argues that the aretaic face of responsibility has ethical depth since such appraisals "implicate one's practical identity."[119] As Watson puts it:

---

[118] Watson (2004) p. 271.

[119] Ibid. The claim that attributability has "ethical depth" is offered in response to a criticism by Wolf (1990) that aretaic assessments are akin to "mere grading" responses and thus aren't deep or significant.

> Aretaic evaluations thus differ significantly from other forms of appraisal. If I dance clumsily, it is inescapably true of me that I was (on that occasion) a clumsy dancer. But if what I do flows from my values and ends, there is a stronger sense in which my activities are inescapably my own: I am committed to them. As declarations of my adopted ends, they express what I'm about, my identity as an agent. They can be evaluated in distinctive ways (not just as welcome or unwelcome) because they themselves are exercises of my evaluative capacities.[120]

Consider a case in which a person does something, but the action doesn't express or disclose "what they're about", i.e. it doesn't reveal anything about their identity as an agent. Say that the person in DRUNK DRIVER is a very conscientious and considerate person; they are not the kind of person who would knowingly engage in behavior that might endanger themselves or other people. In such a case, the erratic driving isn't *attributable* to the agent; it doesn't express or disclose their true commitments or values. Once again, this stands in sharp contrast to the malicious vase smasher. In the latter case, the action *is* attributable to the person since it is an expression or manifestation of their (callous) system of values and commitments.

There are many vexing questions about the relationship between moral accountability (especially blame) and attributability. According to a certain cluster of views that are sometimes called "deep self" or "self-disclosure" views, an action's being *attributable* to an agent is a necessary condition for the appropriateness of moral blame[121]. In particular, such views hold that, in order to be an apt target of blame, the action must somehow reveal or disclose some morally objectionable aspect or feature of the person's "deep" self, e.g. that the person fails to

[120] Watson (2004). pp. 270-71.

[121] Historical antecedents of these contemporary views include Frankfurt (2003) and Watson (2003). Wolf (1990) calls such views "real self" views. Contemporary versions include Sripada (2016) and Shoemaker (2015). Reis-Dennis (2017) argues against "deep self" views by arguing that a person can still be subjected to moral blame (including reactive attitudes) even when their conduct discloses or reveals *nothing* about their "deep self". Reis-Dennis includes Angela Smith's "answerability" view as a version of the "deep self" view. However, I think it's better to keep Smith's view separate. I will discuss Smiths' view momentarily.

accord others a sufficient level of regard, or that the "quality of their will" is somehow deficient or substandard[122]. Once again, such views take this kind of self-disclosure to be a necessary condition for *moral* accountability[123]. I won't take a stand on the question of whether or not attributability is a necessary condition for moral accountability. However, I would like to note that, as we've seen, when it comes to various *non-moral* forms of accountability (e.g. legal) attributability is not necessary.

There is one further notion that I would like to isolate; one that will also be relevant for the subsequent discussion regarding epistemic accountability. This is the notion of *answerability*. Angela Smith has raised some doubts concerning Watson's distinction between responsibility as "accountability" and responsibility as "attributability"[124]. According to Smith, it's not necessary that we multiply senses of agential "responsibility"; we can provide a unified account rooted in the notion of "answerability". While I do not intend to enter this dispute here, it will be helpful to

---

[122] This relates to a central "commonplace" that Peter Strawson sought to remind us of in his essay "Freedom and Resentment":

> The central commonplace that I want to insist on is the very great importance that we attach to the attitudes and intentions towards us of other human beings, *and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions*. (p. 62, emphasis added).

[123] Importantly, this is taken by many as a necessary but *not* sufficient condition for moral accountability. This point is sometimes made by considering cases involving psychopathy. Some argue that taking up a blame response to a psychopath implies that the psychopath has certain capacities, e.g. an ability to recognize the interests of others as making a valid claim on him. However, given that psychopaths lack this capacity, taking up a blame response towards him is perhaps inappropriate. Nevertheless, we can still think of the psychopath as *callous* or *cold*; his actions and attitudes are still expressive of his agency, and he is still subject to certain forms of evaluation. Here, arguably, is a case involving attributability *without* moral accountability (there are, of course, *non-moral* forms of accountability that will be called for in such cases). For arguments against this kind of view on psychopaths see: Smith (2015), (2019).

[124] Smith (2012), (2015).

get the basic idea behind Smith's answerability account on the table[125]. Smith's notion of "answerability" is meant capture the most basic condition of agential responsibility. This notion is subtly different from "attributability". First, as we will see, according to Smith, "answerability" is connected with our *rational* capacities; specifically, our ability to justify our attitudes and actions by citing reasons in support of them. One might deny that "attributability" is tied to reasons or rationality in this way. Second, on Smith's view, we can be "answerable" for things that that aren't necessarily reflective or expressive of our *deep* selves, where by "deep" I mean our *core* commitments and values. For instance, according to Smith, we are, in principle, answerable for rather mundane actions and attitudes (e.g. brushing our teeth or walking to the grocery store).

According to Smith, to say that an agent is responsible for something in the "answerability" sense,

> is to say that that agent is an appropriate target, in principle, of requests for justification regarding that thing and that she is eligible, in principle, for a variety of moral responses depending upon how well or poorly she meets this justificatory request.[126]

Thus, according to Smith, the most basic condition for agential responsibility – whether for actions or attitudes – has to do with rational justification. We are "answerable" for things when it would be appropriate, in principle, to ask us to justify those things by citing reasons. This is what makes us *eligible* for a variety of further reactions (blame, reproach, gratitude, etc.) depending on the quality of our response to this kind of justificatory request. On Smith's view, then, what is significant when it comes to this most basic condition for agential responsibility is not voluntary

[125] Scanlon (1998) also puts forward an answerability account of responsibility.
[126] Smith (2015) p. 103.

choice or control, but whether or not the attitude or the action is related to a person's underlying judgments:

> What matters, on this account, is whether an action or attitude is normatively connected to a person's underlying judgments in such a way that she can, in principle, be called upon to defend it with reasons and to acknowledge fault if an adequate defense cannot be provided. Bodily movements and mental states that are not even in principle answerable to a person's judgment are therefore not the sorts of things for which we are responsible, on this account; but we are responsible for most of our desires, emotions, beliefs, and other attitudes, despite the fact that they do not generally arise from conscious choice or decision and are not normally under our immediate voluntary control.[127]

It will be important to keep these conceptual distinctions – accountability, attributability, and answerability – in mind as we move forward with the discussion.

3. Epistemic Accountability

As I have noted, our accountability practices (whether moral or non-moral) are complex and varied. The question I would like to consider now is the following: Do we have a practice of *epistemic* accountability? Recall that the task at this stage is a descriptive/interpretive task vis-à-vis our existing social practices. I will turn to the question of whether or not our social practices are *legitimate* in the next chapter. For instance, we might find that, as a matter of fact, we hold each other accountable qua believers in ways that are unfair or unjustified[128]. However, before asking questions about the legitimacy of the practice, we must first try to provide an accurate description of the practice itself.

I'd like to begin by noting that holding a person accountable, in general, is going to involve some kind of response in reaction to their thinking or behavior. When it comes to

[127] Smith (2008) p. 370.
[128] Of particular importance will be the phenomenon of "epistemic injustice". I will return to this issue in the next chapter.

epistemic accountability, we'll be looking for a form of response that can be taken up in reaction to a person's violation of a norm of a certain kind. We can avail ourselves of the norms that I earlier described as our "starting materials", e.g. norms which prohibit things like bald inconsistencies, wishful thinking, and believing on the basis of insufficient evidence. Recall that, such norms provide us with a pre-philosophical starting point when it comes to the "epistemic". In availing myself of them, I do not mean to beg any questions; they can be developed in various theoretical directions, and they are amenable to various ways of understanding the "epistemic".

What we're looking for, then, is a response that can be taken up on the part of others in reaction to an individual's violation of a norm of the relevant kind. It's important to keep in mind that not just *any* response will do. Recall the "cool" or stoic observer of the malicious vase smasher that I mentioned earlier. This person recognized that a bad act was carried out and that it issued from an objectionable character. Insofar as the stoic observer "recognized" that a bad act occurred, we might say that he responded to the action (in a rather weak sense of "responded"). However, given that he didn't take up any *further* responses (e.g. he wasn't emotionally exercised by the act, he didn't protest, he didn't demand an apology, etc.), then he didn't count as holding the vase smasher *accountable*. Rather, he was engaged in mere assessment vis-à-vis a standard. Thus, in order to count as an *accountability* response, the response in question must go beyond mere assessment vis-à-vis a standard. Also keep in mind that, while *moral* accountability is an important form of accountability, it is not the *only* form.

We should begin by noting that there are *many* different ways that we can react to a person's violation of a norm of the relevant kind. Here are just a few examples:

1. A teacher responds to a student's fallacious line of reasoning in a paper by giving the student a bad grade.

2. An employee responds with indignation upon learning that her employer believes that women are unlikely to be successful employees. (The employer often hires women, but only to avoid censure).

3. One friend gently rebukes another after learning that he believes things about the personal lives of celebrities based on what he reads in tabloids.

4. One co-worker becomes angered with another after being assured of a "hot stock tip" that's based on flimsy evidence.

These examples should suffice to give us a sense for the wide variety of reactions that can accompany failures that have an epistemic dimension, as well as the broader social contexts within which these responses make sense. However, in each of the cases listed above, it seems that the response is somehow sensitive to considerations *external* to the fact that a person has violated a norm of the relevant kind. In other words, it doesn't seem like any of 1-4 provide us with a response that constitutes a *distinctly epistemic* form of accountability.

Consider, for instance, the response in case 1. The giving of the bad grade in this case is clearly tied up with the institutionalized teacher-student relationship. The response is thus not taken up in reaction to a person's violation of a norm of the relevant kind, *considered as such.*[129] Rather, it is a form of response that is embedded in a more local practice. The locality of the practice means that the form of response will be constrained and modulated by various *other* factors that are unrelated to the question of how well an individual performs vis-à-vis norms of the relevant kind, considered as such. For instance, it's only appropriate to give grades (whether

[129] "Considered as such" is a convenient stand-in at this point in the discussion. As I've mentioned earlier, my hope is that answering the justificatory question vis-à-vis epistemic norms in general will help us better understand the class of "epistemic" norms as a whole, e.g. perhaps it will reveal something about the proper place or purpose of these norms.

good or bad) *to work that has been assigned.* However, an individual can clearly still perform well or poorly vis-à-vis norms of the relevant kind without being given a particular school assignment (not to mention without being a student at the particular institution in question).

Similar points can be made regarding case 2. In addition to being unsubstantiated by the evidence, the employer's belief clearly has *moral* significance[130]. In maintaining his belief about female employees, the employer is failing to *relate to others* in the way that he ought to.[131] To put it in language I introduced earlier, it seems that the employer's belief expresses or discloses some morally objectionable feature of his underlying commitments and values; he fails to accord others a sufficient level of regard. Moreover, it seems that the indignation response taken up by the employee is tracking *this* aspect of the employer's belief rather than the fact that the belief is unsupported by the evidence. For instance, it doesn't seem that the employee would respond with indignation if the evidentially unsupported belief was about some trivial or mundane issue (e.g. imagine an instance where the employer has evidentially unsupported beliefs regarding the biographical details of some obscure 19th century American artist). If this is right, then it seems that the employee is holding the employer *morally*

[130] While I do not have the space to argue for it here, it seems to me that the immorality of the employer's belief isn't simply *reducible* to the mere fact that it's based on insufficient evidence. Rather, it seems to me that there is a *distinctly moral* harm involved in this case, i.e. one arising from the very nature of the belief itself, yet which is unrelated to considerations of evidence. This connects with larger questions concerning the interaction between moral and epistemic demands vis-à-vis our doxastic attitudes. I don't have the space to adequately explore these questions here.

[131] This point is controversial. Some deny that our beliefs are morally evaluable in this way. However, note that Smith's "answerability" view gives us the resources to accommodate our practices of holding each other morally responsible for our attitudes. For further support for the claim that our attitudes are open to moral assessment, see: Basu (2019a), (2019b), Marušić & White (2018).

accountable for his attitude.

However, as we noted earlier, moral accountability responses are sensitive to a host of considerations that appear to be unrelated to the question of how well an individual performs vis-à-vis norms of the relevant kind, considered as such. Recall the issue mentioned earlier having to do with "standing" to engage in moral blame. Sometimes a blame response (e.g. resentment or indignation) can be inappropriate or unjustified due to *complicity*, *hypocrisy*, or *meddling*. Considerations such as these do not seem relevant when it comes to how well a person performs vis-à-vis norms of the relevant kind. Imagine, for instance, a person who would somehow be guilty of hypocrisy if they were to morally blame the employer for his sexist attitudes. Perhaps blame responses such as indignation are now off the table for this person. But is there not some *other* (distinctly *epistemic*) form of accountability which still *is* on the table? After all, the employer's belief is unsupported by the evidence, so it seems to violate a norm of the relevant kind. What we're looking for, once again, is an accountability response that is specifically tied to an individual's performance vis-à-vis norms of this kind. It doesn't appear that cases 1 and 2 have delivered that.

Similar points hold for cases 3 and 4. Regarding case 3, gentle rebukes are also sensitive to a host of ethical considerations that seem to be unrelated to the question of whether or not a person performs well vis-à-vis norms of the relevant kind. For instance, it might be acceptable to gently rebuke a *close friend* for believing some tabloid gossip. However, it would be entirely inappropriate for me to reproach a *complete stranger* who I happen to see reading a tabloid magazine on the bus. Once again, these ethical constraints on the appropriateness of various forms of interpersonal interaction do not seem to be connected in the right way with the question

of how well a person performs vis-à-vis norms of the relevant kind. This issue obtains in our last case as well. In case 4 the anger response seems to be sensitive to the fact that the co-worker was *assured* by their colleague of the "hot stock tip". Being the recipient of an assurance like this – not unlike being the recipient of a promise – generates various expectations and entitlements. Social exchanges such as these could even be considered (crude and informal) ways of *contracting* with one another; they involve various (perhaps unspoken) agreements and expectations that come into play only when the parties to the exchange somehow consent or choose to partake. The anger response in case 4 seems to be tied to the fact that the two co-workers engaged in this kind of exchange. For instance, it doesn't seem that some *third party* (say, another co-worker in the office who was witness to the initial exchange in virtue of occupying an adjacent cubicle) would be entitled to an anger response[132]. But now we've run into the same issue once more; it seems that there will be various forms of response in this case that have appropriateness conditions which aren't connected in the right way to the central question of how well a person performs vis-à-vis norms of the relevant kind, considered as such. The anger response taken up by the recipient of the assurance seems to be a response like this. Thus, we have failed to isolate a *distinctly epistemic* form of response; one which is taken up in reaction to an individual's performance vis-à-vis norms of the relevant kind, considered as such.

We can utilize the forgoing discussion to formulate three conditions that a form of response R would have to meet in order to be counted as a distinctly epistemic form of

[132] Even if an anger response *is* called for on the part of third parties, it doesn't seem like a *distinctly epistemic* form of response (consider, once again, examples involving evidentially unsupported beliefs regarding trivial or insignificant subject matters).

accountability. It's worth reiterating that, at this stage, I'm not assuming that *there is* a response that meets these conditions; one might take the above as evidence against the existence of a distinctly epistemic form of accountability. Nevertheless, if one exists, it seems that it would have to meet the following three conditions:

> A form of response R taken up vis-à-vis S constitutes an existing way of holding S epistemically accountable iff:[133]
>
> (i) R a part of our actual social practices; and
>
> (ii) R is a *distinctly epistemic* form of response, i.e. it is taken up in reaction to S's performance vis-à-vis norms of the relevant kind, considered as such; and
>
> (iii) Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard.

Condition (i) is straightforward; the relevant form of response has to actually be a part of our existing social practices. Condition (ii) says that the response has to be *distinctly epistemic*. This was the condition that failed to obtain in cases 1-4 above. In all of those cases, a form of accountability was perhaps on display. However, as we saw, the accountability responses were all sensitive to considerations that were not directly relevant to the question of how well a person performs vis-à-vis norms of the relevant kind, considered as such. Condition (iii) says that the response has to somehow go beyond mere assessment vis-à-vis a standard. This condition rules out the possibility of an epistemic analogue of the kind of "detached" form of assessment that our stoic observer of the vase smasher engaged in.

Is there a form of response that meets these three conditions? I believe that there is. To begin, it will be helpful to remind ourselves of the various ways in which our doxastic lives are

[133] The inclusion of 'existing' is perhaps awkward, but it is necessary; there might be forms of responses that are not a part of our existing practices but that count as accountability responses. If there are such responses, then they aren't relevant to this discussion.

intertwined with our wider social and environmental settings, and the extent to which we depend on each other when it comes to the acquisition and dissemination of information. There is a traditional, individualistic, conception of the subject matter of epistemology, one which tends to focus on the mind of a solitary agent reasoning in isolation. This solitary agent might (à la Descartes) be reasoning in isolation from the *external world as whole*. However, we don't even need to go that far in order to maintain a kind of epistemic individualism; one might just conceive of this agent as reasoning in isolation *from other people*.

These individualist conceptions are a far cry from our actual, lived experience as doxastic agents. Consider, for instance, the social or interpersonal dimensions of our doxastic lives[134]. It seems that our beliefs can be shaped in profound ways by our relationships with other people, e.g. one's political beliefs might be influenced by one's parents, one's views about various current events might be influenced by one's "friends" on social media, etc. Relatedly, we *depend on other people* when it comes to acquiring information about various topics and subject matters, and we are sometimes *called upon by others* to act as informants. The ways in which we depend on each other in order to form accurate views about the world around us are so pervasive and deeply interwoven into the fabric of our everyday lives that it can be easy to omit them or forsake their importance. It's not possible to provide an exhaustive catalogue, but let's remind ourselves of just a few key ways in which we depend upon each other in this way. There are, of course, instances involving direct, face-to-face communication between two individuals, e.g.

[134] I'll set aside here important questions concerning the view known as "content externalism". According to this view, features of our objective environments (i.e. the nature of our physical surroundings) can partially determine the contents of our attitudes.

when one person wants to know what time the bus is scheduled to arrive and asks the person standing next to him. This is a paradigm case of what epistemologists call "testimonial" interaction. However, if we broaden the conception of "testimony" ever so slightly, then the extent to which we depend upon others in order to acquire information about the world is simply staggering. As Fricker (2006) notes, "testimony" in this wider sense

> can occur through an extensive range of types of spoken and written means of purportedly factual communication, including telephone calls, e-mails and personal letters, lectures and radio broadcasts, newspapers, textbooks and encyclopaedias, personal diaries, and public records of all kind.[135]

While this list is not exhaustive (we might also include text messages, blog posts, social media posts, etc.), it should hopefully give us a sense for the numerous ways in which we depend upon other people in order to arrive at an accurate view of the world. Thus, we should keep these observations about our social-epistemic circumstance in mind as we proceed.

Return, now, to our question: Is there a form of response R which meets criteria (i)-(iii) above? Consider the following: It seems that we can *vary our level of trust and reliance in people* when it comes to their role as testifiers and informants. For instance, a person might be viewed as more or less credible when it comes to various topics, subject matters, and issues. In certain instances where a person is viewed as less credible vis-à-vis a certain topic or subject matter, *other people* might modify their level of trust in the person when it comes to that topic or subject matter. This kind of trust modification might involve a number of things, e.g. it might involve an unwillingness to take that person's claims at face value when it comes to the particular topic or subject matter, or it might involve a reluctance to rely on the person as a

[135] Fricker (2006) p. 592.

testimonial source of information when it comes to the topic or subject matter.

Let's look at a specific example in order to fill out this idea. Return again to case 3 from above; the case where one co-worker was assured by another of a "hot stock tip" that was based on flimsy evidence. As I mentioned earlier, it seems clear that the recipient of this assurance is entitled to certain reactions given that the assurance was based on flimsy evidence. However, we also noted that this reaction doesn't appear to constitute a *distinctly epistemic* form of accountability; it isn't a reaction to the person's performance vis-à-vis norms of the relevant kind, considered as such. Specifically, we noted that it doesn't seem like some *third party* (say, another co-worker in the office who was witness to the initial exchange in virtue of occupying an adjacent cubicle) would be entitled to an anger response. However, imagine that, after hearing the initial exchange, this third co-worker did some quick research and came to find that the "hot stock tip" was bogus. While an anger response might be uncalled for, it *does* seem perfectly acceptable for this person to *modify her level of trust* vis-à-vis the co-worker who offered the assurance. For instance, she might view this person as an unreliable source regarding investment tips and cease to take his words at face value when it comes to this topic. This recognition of unreliability and subsequent adjustment of trust can come about in many different ways. In the case above, the third co-worker happened to overhear a conversation. This is certainly not the only way in which a recognition of unreliability and subsequent adjustment in trust can come about. For instance, the recipient of the initial assurance might also modify trust in the person in this way (in addition to being angry with the person). Also, one might be informed by a trustworthy and dependable friend (or some other reliable source) that a certain person is unreliable vis-à-vis some topic or subject matter. There are various ways in which we can come

to be informed about the reliability of others vis-à-vis certain topics or subject matters.

As I have noted, the kind of trust modification at issue here can take several shapes. For instance, one might cease to trust the person's claims when it comes to a particular topic or issue. "Claims" can be made in many different ways, e.g. via explicit speech acts in the course of verbal communication (assertions, tellings, etc.), in writing (text messages, emails, etc.), and even through certain gestures or movements. Thus, "claiming" that p is, roughly, presenting p as true through some kind of expressive act (whether verbal or otherwise). It's important to note that "claims" (so understood) are often made with the explicit intention of getting *others* to accept as true the proposition which is presented as true (perhaps just on the basis of one's say-so). Modifying trust in a person can involve placing less confidence in expressed claims like these. For instance, one can cease to take a person's expressed claim that p as good reason to think that p is true. This is thus one way in which the relevant kind trust modification can be manifested. However, one's modification of trust in a person can also show up even if the target of the response hasn't yet explicitly expressed some claim(s) vis-à-vis the relevant topic or subject matter. For instance, one might simply *intend* not to trust a person's claims regarding the topic or subject matter; something which can happen even if the target of the response hasn't yet explicitly expressed a claim regarding the issue. Also, one might simply be less inclined to *seek that person out* as testimonial source of information when it comes to that topic or subject matter.

Notice as well that the modifications in trust that I just described can occur privately. In other words, one doesn't have to *overtly express* these attitudinal adjustments, whether directly

to the person who is the target of the response, or to anyone else.[136] Indeed, there might be good reason not to. Return again to our office workers in the "hot stock tip" example. Say that the person who offered the initial assurance was the C.E.O.'s obnoxious and incompetent son. If the office worker who overheard the initial exchange were to overtly express her attitudinal responses directly to this person, that might lead to various other consequences that this office worker would rather simply avoid, especially given the relative innocuity of the episode. We can imagine, for instance, that the C.E.O. is a giant pushover when it comes to his son's whims and desires, and that his son is an arrogant and entitled hothead who often seeks out personal retribution whenever he feels slighted by someone.[137] Even if this is the nature of the case, the office worker who overheard the initial exchange can still adjust her trust vis-à-vis the C.E.O.'s son in the ways described above.

I offer these observations about our ordinary practices as a preliminary case for the claim that the kind of trust modification described above is a distinctly epistemic form of accountability. As we will see, this case is not complete. However, let's offer the following as a (preliminary, first pass) characterization of epistemic accountability:

Epistemic Accountability: A holds B epistemically accountable iff:

(i) A judges that B violates a norm of the relevant kind; and

(ii) A modifies her level of trust in B in ways X; and

(iii) A makes the modifications in (ii) because of her judgment in (i).

[136] This is not to say a person *can't* express these attitudinal adjustments. It's just to say that overt expressions aren't *necessary*.

[137] If this is the nature of the case, would the recipient of the initial assurance still be entitled to an anger response? I think the recipient would still clearly be entitled to this reaction, although there might be prudential reasons not to express it directly to the C.E.O.'s son. Such is the unfortunate predicament we sometimes find ourselves in.

Recall that, in order to actually count as epistemic accountability, this response must meet the three desiderata introduced earlier. I'll repeat those here:

> A form of response R taken up vis-à-vis S constitutes an existing way of holding S epistemically accountable iff:
>
> (i) R a part of our actual social practices; and
>
> (ii) R is a *distinctly epistemic* form of response, i.e. it is taken up in reaction to S's performance vis-à-vis norms of the relevant kind, considered as such; and
>
> (iii) Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard.

Does the proposed account in **Epistemic Accountability** provide us with a form of response R that meets the three criteria? When it comes to the first condition of **Epistemic Accountability** ("A judges that B violates a norm of the relevant kind") it's important to keep in mind that, at this point in the discussion, I'm simply operating with certain norms that I earlier offered as our "starting materials", i.e. commonsense and ordinary prohibitions on things like bald inconsistencies, wishful thinking, or believing without adequate evidence. The hope is that, by identifying a distinctly epistemic form of accountability, we'll be better positioned to understand the class of "epistemic" norms itself. What this means is that the kind of trust modification at issue in condition (ii) of **Epistemic Accountability** ("A modifies her level of trust in B in ways X") will be very important for the ensuing discussion.

As I described above, the kind of trust modification at issue here will involve things like *ceasing to take a person's claims at face value* (perhaps only when it comes to a certain topic or subject matter), where "claiming" that p is, roughly, presenting p as true through some kind of expressive act, whether verbal or otherwise, often with the explicit intention of getting other to

accept as true the proposition which is presented as true. Recall that one can present p as true in this way through various different means. As I have mentioned, there are explicit speech acts offered in the course of verbal communication, e.g. acts of asserting or telling. Note as well that, when it comes to such acts, one can cease to take them at face value even if one isn't within their *intended* audience (as we saw in the case of the office worker and the "hot stock tip"). There are also written communications of various kind, and even certain gestures or movements. Once again, one can cease to take these at face value even if one isn't within their intended audience. When one ceases to take claims like these at face value, one no longer sees them as good reason to think that the proposition which is *presented* as true actually *is* true. In addition to ceasing to take a person's *expressed* claims at face value, one might also *intend* to not trust a person claims regarding some topic or issue. This can happen even if the target of the response hasn't yet expressed any claim(s) regarding the topic. The form of trust modification at issue might also involve a reluctance or unwillingness to rely on a person as a testimonial source of information, perhaps only when it comes to a certain topic or subject matter. A "testimonial" source of information is a person who is relied upon to convey information through certain acts, e.g. speech acts of telling or asserting. Once again, such expressive acts involve making "claims", in the manner described above.

If we take this as our preliminary characterization of the relevant forms of trust modification at issue in **Epistemic Accountability** (i.e. if we take this as a preliminary way of filling out the relevant "ways X" in "A modifies her level of trust in B in ways X") will **Epistemic Accountability** meet the three criteria for our desired response R? Let's take them one at a time. When it comes to the first condition ("R is a part of our actual social practices") I

think the account is in good shape; we do frequently modify our level of trust in each other in way similar to my characterization above. How about the second condition ("R is a *distinctly epistemic* form of response, i.e. it is taken up in reaction to S's performance vis-à-vis the relevant norms, considered as such")? Here it seems like the account on offer represents a promising way of fleshing out what a "distinctly epistemic" form of accountability might look like. For instance, it doesn't look like the account on offer confronts the challenges associated with the forms of response that we saw at play in cases 1-4 from earlier. Recall those cases:

> 1. A teacher responds to a student's fallacious line of reasoning in a paper by giving the student a bad grade.
>
> 2. An employee responds with indignation upon learning that her employer believes that women are unlikely to be successful employees. (The employer often hires women, but only to avoid censure).
>
> 3. One friend gently rebukes another after learning that he believes things about the personal lives of celebrities based on what he reads in tabloids.
>
> 4. One co-worker becomes angered with another after being assured of a "hot stock tip" that's based on flimsy evidence.

As we have seen, in each of these cases it seems as though the response which is taken up vis-à-vis the relevant individual is sensitive to considerations unrelated to the question of how well that individual performs vis-à-vis norms of the relevant kind, considered as such. However, it does not look like the proposal in **Epistemic Accountability** confronts this issue. For instance, in contrast to the grading response offered in case 1, the target of the response captured in **Epistemic Accountability** doesn't have to be one's *student*; nor do they have to be a student at all. Similarly, in contrast to the indignation response described in case 2, the target of the response captured in **Epistemic Accountability** doesn't *also* have to be the appropriate target of

moral accountability. Recall again the example of the person who has lost "standing" to blame others for sexist attitudes. Even if this person has lost standing to blame the employer in case 2 for his sexist attitudes, he can certainly still modify his level of trust in him in the manner described above. In other words, even if one has lost standing to morally blame someone for their attitudes, they can still retain standing to hold them epistemically accountable in the manner described above. Similarly, in contrast to the gentle rebuking in case 3, the response captured in **Epistemic Accountability** isn't sensitive to a host of ethical considerations regarding the appropriateness of overt acts of reproach or criticism. And finally, in contrast to the anger response in case 4, the response captured in **Epistemic Accountability** isn't beholden to considerations involved in certain forms of complex social interaction, e.g the reciprocal expectations and associated entitlements that arise in the context of acts of assuring, promising, or telling. Rather, as we saw in the case of the third-party office worker, someone can take up the responses captured in **Epistemic Accountability** vis-à-vis a person even if they have *not* engaged with this person in a cooperative social act of this kind.

Thus, it seems that the proposed account in **Epistemic Accountability** does well when it comes to meeting the first two conditions for our desired response R. But what about the third condition ("Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard")? This third condition is undeniably the hardest of the three conditions for **Epistemic Accountability** to meet. Nevertheless, I think that this third condition is met as well. I will spend the remainder of the present chapter arguing that **Epistemic Accountability** meets this third condition.

The first thing to note is that the second condition in **Epistemic Accountability** ("A modifies her level of trust in B in ways X") is what is supposed to take the response captured in **Epistemic Accountability** beyond mere assessment vis-à-vis a standard. Recall, once again, our "cool" or stoic observer of the malicious vase smasher. That person observed that a bad act was carried out, and that it issued from an objectionable character, but no *further* responses were taken (e.g. the stoic observer wasn't emotionally exercised, he didn't protest, he didn't' demand an apology, etc.). The "further" response which is supposed to distinguish the response captured in **Epistemic Accountability** from mere assessment vis-à-vis a standard is the fact that A *actually modifies* her level of trust in B in ways X. It seems that an individual A could correctly judge that another person B violates a norm of the relevant kind *without* actually modifying her trust in B in ways X. For instance, A might correctly judge that B violates a norm of the relevant kind, but then make no further adjustments when it comes to her trusting attitudes vis-à-vis B; maybe A will never see B again, or maybe A just forgot, or perhaps she just didn't think that this particular norm violation was all that important, or maybe she wants to give B a second chance, etc. Thus, it's the *actual modification* in trust, taken up by A vis-à-vis B, which is supposed to distinguish the response captured in **Epistemic Accountability** from mere assessment vis-à-vis a standard.

Nevertheless, one might remain doubtful that the response captured in **Epistemic Accountability** actually *does* go beyond mere assessment vis-à-vis a standard. I'd like to mention one way of developing this point before going on to discuss what I think is the more challenging form of this objection. The first way of developing the above point appeals to the fact that we can modify our level of trust in inanimate objects and artefacts. For instance,

imagine finding out that a particular thermometer is faulty and that it delivers inaccurate readings. Upon learning this, a person can cease to trust the readings provided by the thermometer. Hasn't this person responded to the thermometer in the same way that the target of the response captured in **Epistemic Accountability** is responded to? And if so, how can the form of response captured in this account amount to an *accountability* response? Surely we don't hold *thermometers* accountable!

The reply here is fairly straightforward: The way that we respond to faulty thermometers and other unreliable objects and artefacts is *not* identical to the form of response captured in **Epistemic Accountability**. The key clause in **Epistemic Accountability** is, once again, clause (ii): "A modifies her level of trust in B in ways X". In particular, the "ways X" are crucial for answering the present worry. The relevant trust modifications captured in **Epistemic Accountability** involve ceasing to take *a person's claims* at face value (perhaps only when it comes to certain topics or subject matters), or, alternatively, a reluctance or unwillingness to rely on a person as a *testimonial source of information* (where a "testimonial" source of information is someone who conveys information to another by making claims). "Claiming" that p, once again, involves presenting p as true through some kind of expressive act (whether verbal or otherwise), often with the explicit intention of getting others to accept as true the proposition which is presented as true. Thermometers do not make "claims" at all, so understood; they don't engage in expressive acts, and they don't have intentions. One might say, for instance, "The thermometer reads 98.6" and infer from this that they don't have a fever. However, statements like "The thermometer *told me* that my temperature is 98.6" or "The thermometer *assured me*

that my temperature is 98.6" can't be taken literally; we have to read them as colloquial ways of expressing what's stated in the first claim, or as jokes. While a *doctor* might "tell" or "assure" you of such things, a thermometer's reading is only a mere sign or indication. One way to capture this idea is in terms of a distinction originally put forward by Edward Craig (1990) between "informants" and states of affairs with evidential import:

> among the various sources of information there are on the one hand informants who give information; and on the other there are states of affairs, some of which involve states of human beings and their behavior, which have evidential value: information can be gleaned from them. Roughly, the distinction is that between a person's telling me something and my being able to tell something from observation of him.[138]

When it comes to the form of response captured in **Epistemic Accountability,** only "informants" (in Craig's sense) are the appropriate targets. "Informants" are individual's who are capable of making "claims". As a result, they can participate in various forms of cooperative communication, e.g. they can "tell" us things where this involves more than merely manifesting behavior that has evidential import. Thermometers and other inanimate objects and artefacts, not being informants, are thus not responded to in the ways captured in **Epistemic Accountability**[139].

---

[138] Craig (1990) p. 35.

[139] Some authors have argued that an adequate account of the epistemology of testimony must somehow take account of Craig's distinction between "informants" and states of affairs with evidential import. In particular, a cluster of views that are sometimes called "assurance" views hold that the testimonial exchange is irreducibly *second-personal* in nature; it essentially involves two agents relating to one another *as* agents rather than as objects of management or prediction. Proponents of such views hold that these second-personal elements have an important role to play when it comes to understanding the epistemology of testimony. This view is often traced back to Ross (1986). Recent proponents include: Faulkner (2011), Hinchman (2005, 2014), McMyler (2011), and Moran (2006, 2018). Arguments against this view have been given by: Kornblith (2022), Lackey (2008), and Schmitt (2010). While I find Craig's distinction helpful, and I agree that testimonial encounters will often involve second-personal elements, here I remain neutral on the question of the relevance of such elements for understanding the epistemology of testimony.

Setting this first objection aside, let's now discuss the more challenging version of this worry. One could grant the above response to the thermometer concern but still persist: Even if the kinds of trust modification captured in **Epistemic Accountability** are *only* taken up vis-à-vis "informants", that *alone* doesn't establish that **Epistemic Accountability** meets the third condition for our desired response R ("Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard"). We still need some positive account which establishes that the response captured in **Epistemic Accountability** meets this third condition. In the next section I turn to this question.

4. The Significance of Epistemic Accountability

Why does the form of response captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard? Amongst contemporary authors, Cameron Boult has probably devoted the most attention to this question[140]. Boult has developed an account of what he calls "epistemic blame" that is influenced by T.M. Scanlon's relationship-based account of moral blame[141]. According to Boult, "epistemic blame" is *also* a kind of relationship-modification. Boult associates his view with a certain family of views (of which Scanlon's view is a part) which seek to understand moral blame in terms of interpersonal relationships.[142]. Indeed, this is how Boult seeks to understand the significance of the responses captured in **Epistemic Accountability**. The rough idea is as follows: The reason why the responses captured in

[140] Boult (2021a), (2021b), (2021c), and (2020). Jessica Brown has also devoted attention to this question (Brown (2019), (2020)). Brown develops an account of epistemic blame based on Sher's desire-based view of moral blame (Sher (2006)). Here I will focus on Boult's view.
[141] Scanlon (2008).
[142] Boult associates his view with Strawson (1962/2003), Scanlon (2008), (2013), and Wallace (2011).

Epistemic Accountability go beyond mere assessment vis-à-vis a standard is that, when one takes up these responses towards a person, one is thereby *modifying one's relationship* with that person, and modifying one's relationship with a person is something that goes beyond mere assessment vis-à-vis a standard; it is an *active* form of response, although it need not involve reactive emotions like resentment and indignation.

Importantly, Boult argues that the responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard when *and only when* they constitute a modification of one's relationship with another person *amounting to blame*. While I think Boult's appeal to a relationship-based framework represents a promising way of understanding *why* the responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard, I do not think his view adequately answers the question of *when* such responses go beyond mere assessment vis-à-vis a standard. Specifically, I do not think Boult is correct in claiming that the responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard *only* when they constitute a modification of one's relationship with another amounting to blame. Let's get this commitment of Boult's on the table. I'll call it "Boult's Claim":

> **Boult's Claim:**
> The responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard *only* when such responses constitute a modification of one's relationship with another amounting to blame.

I am willing to grant Boult the claim that there *are* instances in which a person takes up the responses captured in **Epistemic Accountability** in a way that constitutes a (perhaps *distinctly epistemic*) form of blame. One could push back here, but I won't. My main concern is to argue that **Boult's Claim** is false. **Boult's Claim** is false, I will argue, given that *there are* instances in

which someone takes up the responses captured in **Epistemic Accountability** towards a person where this goes beyond mere assessment vis-à-vis a standard but where this does *not* amount to blaming the person *in any way*. My view is that the relationship-based account gives us the resources to see *why* such instances involve more than mere assessment vis-à-vis a standard. However, as I've noted, the notion of blame is inapplicable in such cases. Thus, I will advance a relationship-based account of the significance of **Epistemic Accountability** which is compatible with, but which does not depend upon, the existence of "epistemic blame".

Before proceeding, let me say more about the basic framework that Boult is operating with. As I've mentioned, Boult draw's influence from T.M. Scanlon's relationship-based account of moral blame. According to Scanlon, our relationships are constituted by certain intentions, expectations, and attitudes we have towards one another. Scanlon distinguishes between relationships as *normative ideals* from *token instances* of relationships. The latter are actual relationships that approximate to a better or worse degree the normative ideal. The "normative ideal" for a relationship – whether friendships, romantic relationships, or otherwise – consists in those intentions, expectations, and attitudes which would comprise a *good* relationship.

According to this framework, a *judgment* of blameworthiness is a judgment that someone with whom you stand in a certain relationship has intentions, expectations, and attitudes that somehow fall short of the "normative ideal" for relationships of that kind. This, however, is not enough to *blame* someone. To actually *blame* a person, one has to *modify one's own* intentions, expectations, and attitudes vis-à-vis the person in a way made fitting by the initial judgment of blameworthiness.

Since I will not be concerned with the question of whether or not Scanlon's approach

represents a satisfactory way of understanding moral blame, I won't spend any more time considering his view. My interest is whether or not the relationship-based view can help us understand the significance of the responses captured in **Epistemic Accountability**. As I've mentioned, Boult utilizes the relationship-based framework in order to develop an account of what he calls "epistemic blame". According to Boult's view, judgments of "epistemic blameworthiness" are judgments that someone has attitudes that impair their *epistemic* relationships with others. To "epistemically blame" someone, then, is to modify one's own attitudes vis-à-vis the person in a way made fitting by this judgment of blameworthiness. But what are "epistemic relationships"? Here's Boult:

> On my approach, epistemic relationships are sets of intentions, expectations, and attitudes people have towards one another that are oriented towards their epistemic agency in distinctive ways. A promising way of approaching this idea is to focus on the fact that, whether we are aware of it or not, most of us have intentions and expectations of one another simply in virtue of the fact that we are epistemic agents in an epistemic community. This is another way of putting the platitude that we rely on one another in a generic way as sources and distributors of information. We rely on one another for such basic things as finding out where the nearest grocery store is, or when the next bus arrives.[143]

According to Boult, members of an epistemic community stand in a relationship of mutual epistemic trust; he calls this "the general epistemic relationship". Thus, what it is to "epistemically blame" somebody, according to Boult, is to modify one's intentions and expectations vis-à-vis the person (e.g. by revising one's intention to trust their words when it comes to a certain topic or issue) in a way made fitting by a judgment *that the person has done something which impairs their epistemic relationship with others*. Impairments of such

[143] Boult (2020) p. 524.

relationships consist in falling short of the "normative ideal" for relationships of that kind. But what is the "normative ideal" for epistemic relationships, and what kind of conduct constitutes "falling short" of that ideal? According to Boult the "normative ideal" for the general epistemic relationship,

> specifies that each member intends to epistemically trust another unless one has good reason not to. And each member expects others to actually *be* epistemically trustworthy in this way, as well as that they will do the same of them.[144]

What are the ways in which people can "fall short" of this normative ideal? Boult mentions instances involving *dogmatism*, *wishful thinking*, *hasty reasoning*, and certain kinds of *biased cognition*. What these failing have in common, according to Boult,

> is that they are examples of intellectual conduct that tends to give others good reason to suspend their presumption of epistemic trust in would-be targets of epistemic blame, at least within some restricted domain, or on some specific matter.[145]

It is crucial for Boult's account that the individual who "epistemically blames" another has a certain kind of judgment, viz. a judgement that the target of the response has done something to impair the general epistemic relationship. This is a *necessary condition* for epistemic blame, on Boult's account. Moreover, as the above perhaps makes clear, Boult has a certain view of what "impairing the general epistemic relationship" must come to. According to Boult,

> epistemic blamers tend to judge that someone has been *intellectually irresponsible*, or *intellectually vicious*, or *reckless*, or just plain *"stupid"*. Those are the sorts of things I take the notion of a judgment of general epistemic relationship impairment to unify. *Only when an agent modifies their epistemic expectations in a way made fitting by this sort of judgment do they count as epistemically blaming others.*[146]

The last italicized sentence will be very important for my argument against **Boult's Claim.**

---

[144] Ibid. p. 525.
[145] Boult (2021b) p. 11.
[146] Ibid. p. 12.(emphasis added to the last sentence).

Recall that, according to **Boult's Claim**, the responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard *only* when such responses constitute a modification of one's relationship with another amounting to blame. As we've just seen, according to Boult's account of "epistemic blame", A epistemically blames B *only if* A makes a certain kind a judgment, viz. that B somehow impairs the general epistemic relationship by being *intellectually irresponsible, intellectually vicious, reckless* or "just play pain '*stupid*'". It is thus an implication of Boult's view that the responses captured in **Epistemic Accountability** go beyond mere assessment vis-à-vis a standard only when the person who takes up those responses has a judgment of this latter kind.

In order to see the error in **Boult's Claim**, it will be helpful to return to the discussion from earlier concerning the threefold distinction between accountability, attributability, and answerability. One way to state my objection against **Boult's Claim** is as follows: In aligning his view of "epistemic blame" so closely with certain accounts of moral blame based in interpersonal relationships, Boult has ended up with a kind of "*doxastic* self-disclosure" view. However, a view like this threatens to distort our understanding of the significance of the responses captured in **Epistemic Accountability**. Recall that, according to "self-disclosure" views, an action's being "attributable" to an agent is a necessary condition for the appropriateness of moral blame. In particular, such views hold that, in order to be an apt target of blame, the action must somehow express or disclose some morally objectionable aspect or feature of the person's "true" or "deep" self, e.g. that the person fails to accord others a sufficient level of regard, or that the "quality of their will" is somehow deficient or substandard. Once

again, such views take this kind of self-disclosure to be a necessary condition for *moral* accountability.

Boult's view implies that, in order to "epistemically blame" a person, their conduct must somehow express or disclose some objectionable feature of what we might call their "*doxastic* self". The existence of a "doxastic self" doesn't require the truth of doxastic voluntarism, i.e. it doesn't require that we have the ability to believe at will. Even if doxastic voluntarism is false, our beliefs can still be expressive or disclosive of who we are in various ways. For instance, an individual might be *hasty* or *lazy* in believing a certain claim; someone might be *dogmatic* or *biased* in maintaining a certain belief; one might be *diligent, careful*, or *charitable* when maintaining their beliefs; and so on. I take the italicized words to be apt ways of characterizing an individual qua believer, i.e. these are ways of describing one's "doxastic self".

Moreover, as is clear from these examples, there seem to be better and worse "doxastic selves". A careful and diligent doxastic agent is one who will seek out and attend to relevant evidence, strive to conform their beliefs to this evidence, try to block the influence of various biases in their thinking, etc. We commonly take these to be markers of a praiseworthy doxastic self. Alternatively, as noted by Boult, we commonly admonish things like hasty reasoning, biased thinking, and dogmatism. An important complication to be noted here is that there might be different *ways* in which a "doxastic self" can be assessed as good or bad. For instance, a doxastic self might be assessed from the *moral* point of view. We might ask ourselves whether or not, *qua believers*, we relate to other people in a favorable way[147]. Also, a doxastic self might be

[147] This idea is taken up and defended in Marušić & White (2018). According to Marušić & White, "one wrongs another person when one's *beliefs and judgments* fall short of the regard the other is entitled to expect from one." (p. 101, emphasis added).

assessed from a strictly intellectual or epistemic point of view[148]. Here I'll remain neutral on the question of whether or not these different points of view pull us in different directions (e.g. whether or not a good doxastic self from the moral point requires things that are in tension with the demands of the epistemic point of view). Since my focus is on **Epistemic Accountability**, I won't be concerned with assessments of one's doxastic self from the moral point of view.

On Boult's view, then, "epistemic blame" is only appropriate when the person who is being blamed has an *objectionable* doxastic self, e.g. when the person believes in ways that are *biased, hasty, dogmatic*, etc. Given **Boult's Claim**, this means that the responses captured in **Epistemic Accountability** will go beyond mere assessments vis-à-vis a standard only when they are taken up towards a person who has a doxastic self of this kind. As I will now argue, this is false. Consider, for instance, cases involving *non-culpable* ignorance. To make this idea clear, let's consider two different cases involving a person who denies the existence of human-induced climate change. In the first case (**Case 1**) the person in question is culpable, in the second case (**Case 2**) the person in question is not culpable:

> **Case 1**: Tom has an abundance of available evidence which is such that, were he to consider it more carefully, he should be rationally compelled to abandon his persistent denial in the existence of human-induced climate change. In spite of this, Tom maintains his belief that human-induced climate change isn't real.
>
> **Case 2**: Bill lives in a community run by a group of individuals who have managed to shut-off access to evidence from the outside world pertaining to the existence of human-induced climate change. Bill is interested in forming true beliefs about this topic, frequently going to the library to research it. Nevertheless, given the evidence available to him, Bill believes that human-induced climate change isn't real.

[148] This is not to deny that "epistemic" assessment of a person's doxastic self might be tied to the social realm in important ways; I'm inclined to think that it is. It is to say that, even if social, epistemic assessment of a person qua believer is distinguishable from moral assessment.

According to Boult's view, Tom is an apt target of epistemic blame; he manifests a dogmatic and irresponsible "doxastic self" insofar as he persists in believing something which is not well supported by the evidence available to him. As a result, according to Boult, we can take up the responses captured in **Epistemic Accountability** towards him. Since this is an instance of epistemic blame, Boult would argue that taking up these responses towards Tom goes beyond mere assessment vis-à-vis a standard.

However, when we turn to Bill in **Case 2**, similar points cannot be made. Bill doesn't manifest a dogmatic and irresponsible "doxastic self" in believing as he does. In fact, he is quite careful and diligent in his efforts to form true beliefs about climate change. Nevertheless, he has an abundance of false beliefs pertaining to this topic. Given the necessary condition for epistemic blame discussed above, Boult must say that Bill is not an appropriate target of epistemic blame. However, given **Boult's Claim**, this means that, if we take up the responses captured in **Epistemic Accountability** in response to Bill, this will *not* go beyond mere assessment vis-à-vis a standard. This, I submit, is incorrect; even though Bill is in no way blameworthy for believing as he does, we can still take up the responses captured in **Epistemic Accountability** to him in a way that goes beyond mere assessment vis-à-vis a standard. The nature of the case may lead one to wonder how we might be able to take up these responses to Bill if *we're* not also members of his community. Here I'd like to note two points: First, recall that, as I mentioned earlier, there are many ways for an individual to make "claims". In addition to explicit speech acts such as telling or asserting, one can also make claims through written communication (e.g. through email, letters, text messages, blog posts, social media posts, etc.). Thus, we can imagine that

Bill's community isn't *so* shut off from the outside world that he doesn't have access to certain ways of making "claims" that can reach an outside audience. As a result, individuals outside of Bill's insulated community can cease to take Bill's claims at face value when it comes to the topic of climate change. Second, we can also imagine a scenario in which Bill *leaves* his insulated community and enters the outside world. At some time *after* Bill leaves his community but *before* he confronts and properly digests evidence which shows that his existing views on climate change are false (something which may take some time), individuals can still take up the responses captured in **Epistemic Accountability** towards Bill. Thus, the responses captured in **Epistemic Accountability** can be taken up towards Bill.

But the question remains: Why think that, in the case of Bill, these responses *do* go beyond mere assessment vis-à-vis standard? Here we can appeal to the relationship-based framework without depending on the existence of "epistemic blame". In responding to Bill in the ways captured in **Epistemic Accountability**, we will be modifying our relationship with him in a certain way. As we saw earlier, we depend on each other in countless ways when it comes to the acquisition and transmission of information. Moreover, the ways in which we depend on each other in this manner seem importantly different from the ways in which we rely on inanimate objects. As we noted earlier, the responses captured in **Epistemic Accountability** are not taken up towards inanimate objects like faulty thermometers. Thus, it seems that Boult's points about the "general epistemic relationship" are apt; we do stand to each other in way that is oriented around the acquisition and dissemination of information. To say that this is a kind of "relationship" is not to say that it is the most *important* relationship in a person's life. However, its status as a relationship imbues it with a level of significance that properly accounts for the

fact that taking up the response captures in **Epistemic Accountability** goes beyond mere assessment vis-à-vis standard.

We can view the accountability responses taken up vis-à-vis Bill as analogous to the responses taken up in certain other instances involving non-moral forms of accountability. Return again to the case of DRUNK DRIVER:

> DRUNK DRIVER: [S]uppose it is necessary for complying with the norm of avoiding reckless driving that your blood alcohol level is below 0.1% by volume. You've had a few drinks that a reliable friend told you had very low alcohol content, and you don't feel drunk, so that you're highly justified in believing that your blood alcohol level is below the required level, and you decide to drive home. But in fact you are over the limit, drive erratically, and are caught by the police. In this case, you may well have an excuse from moral sanctions — in the circumstances, it wouldn't be fair to resent you, say. But you plausibly don't have an excuse from legal sanctions — you should pay your fine. The legal sanction doesn't imply that you're morally bad, or that there's something wrong with your character or your will.[149]

Once again, we can say that the driver in this case is not responsible for his reckless driving in the "attributability" sense of responsibility; he is a careful and conscientious person who wouldn't knowingly put himself or anyone else in danger. Nevertheless, insofar as he engaged in a certain dangerous act (albeit, blamelessly), he is still the appropriate target of legal accountability responses. We might think of Bill in **Case 2** in an analogous fashion. In this case, the response in question is not a sanction imposed by the state. Rather, the response is imposed by other people in one's epistemic community, and it amounts to a modification of one's epistemic relationship with others. Nevertheless, not unlike DRUNK DRIVER, the response can be called for even when the person in question is blameless and has not disclosed or revealed anything objectionable about who they are. It seems perfectly natural to describe the responses in

[149] Kauppinen (2018) pp. 8-9.

DRUNK DRIVER as a form of accountability (albeit, a non-moral form). My suggestion is that we can understand the case of Bill in **Case 2** in an analogous way.

Boult thus seems to get the case of Bill in **Case 2** wrong. Boult is forced to say that, if we do take up responses like those captured in **Epistemic Accountability** in reaction to Bill, this is not something which goes beyond mere assessment vis-à-vis a standard. Rather, Boult must classify it as mere differential reliance, i.e. a response similar to the one that we would take up towards unreliable objects and artefacts. This seems to get this case badly wrong. The kind of response that Boult would have us take up towards Bill is an entirely "detached" or objective response, e.g. one that might be rooted in purely pragmatic considerations related to regulation and management. However, this seems like the wrong way to classify our responses to Bill in **Case 2**. Note the following features of that case: (i) Bill's beliefs about climate change do not express or disclose anything objectionable about his "doxastic self"; (ii) Bill is still "answerable" for his beliefs regarding climate change (i.e. he could still be appropriately asked to justify these beliefs); and (iii) Bill's beliefs regarding climate change are *responsive* to new evidence, i.e. *were* he to acquire strong evidence that his current views were wrong he would revise his beliefs. The "detached" form of response described above doesn't seem apt given (i)-(iii). We might take up a response like this to someone who is, say, incapacitated or somehow impaired. However, in taking up the responses captured in **Epistemic Accountability** towards Bill, we're not doing this. We continue to view Bill as someone who is in full possession of his rational capacities and who is, in a way, performing flawlessly given his situation. We do not, in other words, adopt a stance towards Bill of regulation, management, or prediction. Nevertheless, when taking up the

responses captured in **Epistemic Accountability** towards Bill, we are subjecting him to form of response that goes beyond mere assessment vis-à-vis a standard.

Given the nature of the case, there might be certain responses that Bill is *excused* from. For instance, if there is such a thing as "epistemic blame", then Bill is perhaps excused from that response. However, just as the person in DRUNK DRIVER is excused from certain *moral* forms of accountability but not from *legal* forms of accountability, Bill is not excused from being responded to in the way captured in **Epistemic Accountability.** Also, I of course agree that Tom in **Case 1** can also be subjected to the forms of response captured in **Epistemic Accountability.** Given that Tom *does* manifest an objectionable "doxastic self", these responses might even be more wide-ranging. One might, following Boult, call this an instance of "epistemic blame". Here I remain neutral on that particular question.

Finally, one might think that it's our responsibility to do more than *privately* take up the responses captured in **Epistemic Accountability** in response to Bill in **Case 2**. Imagine again the scenario described above where Bill leaves his insulated community and enters the outside world. In such a case shouldn't we do more than merely *privately* take up the responses captured in **Epistemic Accountability**? Here we return again to issues regarding "standing". In *certain cases*, it might be appropriate (or perhaps even obligatory) for us to present Bill with evidence that he has not yet considered regarding this subject matter. For instance, if we are Bill's teacher, or if Bill has chosen to engage with us in conversation, etc. However, and once again. **Epistemic Accountability** is a form of response that is *distinctly epistemic*, i.e. it is taken up in response to a person's performance vis-à-vis norms of the relevant kind, *considered as such.* Questions of when it would be appropriate for us to try to correct Bill's mistaken beliefs are sensitive to a host

of considerations that are external to the question of how Bill does with response to the relevant norms, considered as such. Thus, even if we don't have "standing" to attempt to correct Bill's mistaken beliefs, we can still take up the responses captured in **Epistemic Accountability** without expressing this directly to Bill.

5. Conclusion

In this chapter I have argued that our social practices of interpersonal criticism and assessment include a distinctly epistemic form of accountability. I appealed to the threefold distinction between accountability, attributability, and answerability in order to frame the discussion. I then formulated three criteria that our desired response R must meet: (i) R a part of our actual social practices; and (ii) R is a distinctly epistemic form of response, i.e. it is taken up in reaction to S's performance vis-à-vis norms of the relevant kind, considered as such; and (iii) Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard. I then proposed **Epistemic Accountability** as a form of response that meets the three conditions. I noted that that hardest of the three conditions for **Epistemic Accountability** to meet is the third condition: "(iii) Taking up R vis-à-vis S goes beyond mere assessment of S vis-à-vis a standard". I argued that **Epistemic Accountability** meets this condition by appealing to the "relationship-based" framework adopted by Cameron Boult. However, I argued that **Boult's Claim** is false; there are instances where the responses captured in **Epistemic Accountability** are taken up towards a person where this does not amount to blame yet it still goes beyond mere assessment vis-à-vis a standard.

Recall that my overarching task is to try to understand evidential normativity. In Chapter

2, I argued that this task is best approached by considering which evidential norms of belief are correct epistemic norms of belief. In particular, I focused on the following norm:

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

The question, then, is whether or not EN is a correct epistemic norm of belief. My suggestion is that we begin with the justificatory question vis-à-vis epistemic norms in general first. My hope is that, by beginning with the justificatory question about epistemic norms in general first, we will be assisted with the content and justificatory questions vis-à-vis evidential norms of belief (including EN). However, I've noted a significant roadblock confronting this approach: there doesn't seem to be any non-arbitrary way of demarcating the class of "epistemic" norms. The way I suggested overcoming this roadblock was by discerning a *distinctly epistemic* form of response that can be taken up in reaction to violations of norms of the relevant kind. The task of the present chapter was to establish that we do, as a matter of fact, have a practice of epistemic accountability. However, as I've noted, this is simply a descriptive/interpretive claim regarding our existing social practices. My hope is that our practice of epistemic accountability will help us answer the justificatory question vis-à-vis epistemic norms in general and thus, by extension, the justificatory and content questions vis-à-vis evidential norms (including EN). The rough idea is that we can provide an answer to the justificatory question vis-à-vis epistemic norms in general like the following: We should accept epistemic norms given that they structure a legitimate social practice that we participate in. However, I have yet said anything about the *legitimacy* of our practice of epistemic accountability. In the next chapter I will turn to this question.

Chapter 5: Epistemic Accountability as a Legitimate Social Practice

1. Introduction

In the previous chapter I argued that we have a practice of epistemic accountability. If my arguments are on the right track, then we have a possible way of specifying the "epistemic" itself; "epistemic" norms would be those that structure a certain kind of social practice. Additionally, we now have a potential answer to the justificatory question vis-à-vis epistemic norms in general: We should accept epistemic norms given that they structure a social practice that we participate in. However, this answer to the justificatory question will only work if the social practice of epistemic accountability is a *legitimate* social practice. In the current chapter I will argue that the social practice of epistemic accountability is indeed a legitimate social practice.

As I've noted, my hope is that, by answering the justificatory question vis-à-vis epistemic norms in general, we'll be able to answer the content and justificatory question vis-à-vis evidential norms (including EN). In other words, this approach will hopefully provide our *reasons for accepting* certain norms. This is what the "justificatory question" asks for when it comes to a certain norm of family of norms; the reason(s) that ground our acceptance of them. The results of the last chapter could potentially help us with this question when it comes to epistemic norms. However, if the social practice of epistemic accountability isn't a *legitimate* social practice, then *even if* we could secure the content of certain norms by looking to the practice, that wouldn't give us a satisfactory answer to the justificatory question vis-à-vis these norms; they would simply be a reflection or codification of an illegitimate social practice.

Thus, in the present chapter I will turn to the question of the legitimacy of our practice of

epistemic accountability. As we saw in the last chapter, the form of response captured in **Epistemic Accountability** can be taken up in response to individuals within our epistemic community. For instance, I argued that in the case of Bill (the person who non-culpably denies the existence of human induced climate change), we can take up the responses captured in **Epistemic Accountability** and thereby subject him to a form of response that goes beyond mere assessment vis-à-vis a standard. One might take this as an accurate characterization of our existing practices, but then *deny* that such a practice gives rise to norms that we should accept. For instance, perhaps we should *reject the practice itself* instead.

I will consider two general strategies for arguing that our practice of epistemic accountability is illegitimate. I will argue that neither of them succeeds in undermining the legitimacy of the practice. First, one might appeal to instances of epistemic injustice as reason for thinking that our practice of epistemic accountability is an illegitimate social practice. "Epistemic injustice" encompasses a wide range of harms that individuals are subjected to in their capacity as knowers and testifiers. A key instance of epistemic injustice is known as "testimonial injustice". This occurs when someone is unjustly accorded a lower level of trust as a result of prejudicial attitudes towards their race, class, sexual orientation, or gender. Given this harmful feature of our social epistemic practices, one might argue that our practice of epistemic accountability isn't legitimate. This worry is particularly pressing for my approach given that I'm trying to derive norms from our existing social practices. Second, one might take issue with our practices when it comes to individuals like Bill, i.e. individuals who are non-culpable yet are nevertheless subject to the responses captured in **Epistemic Accountability**. In particular, one

might think that responding to individuals like this in a way that goes beyond mere assessment vis-à-vis a standard is somehow unfair or unjust. I will argue that neither of these general strategies succeeds in undermining the legitimacy of our practice of epistemic accountability.

Before moving forward, I would like to make explicit a certain methodological assumption that I will be operating with when it comes to practices and their legitimacy. At this point in the discussion, I take myself to have established that we do, in fact, have a practice of epistemic accountability. As I noted above, I will be considering a number of challenges to the legitimacy of this practice. When it comes to practices and their legitimacy, I will assume a "default and challenge" structure. In other words, I will assume the following: An existing social practice X is legitimate if and only if there is no successful challenge to its legitimacy. Note that a "successful challenge" can exist even if it's not explicitly articulated by any person or group of people. What matters is the *existence* of a successful challenge, not the actual articulation and offering of the challenge by some person or group of people. Below I will articulate a number of possible challenges to our practice of epistemic accountability, arguing that none of them succeeds in undermining the legitimacy of the practice. If there aren't any successful challenges to the legitimacy of our practice of epistemic accountability, then the practice is legitimate.

2. Epistemic Injustice

The first issue I would like to consider concerns the phenomenon of epistemic injustice. Consider a case which is offered by Fricker (2007) as a paradigm instance of epistemic injustice. Fricker presents an example from the film *The Talented Mr. Ripley*. This particular scene involves Marge Sherwood and Herbert Greenleaf. Greenleaf is the father of Marge's missing fiancé, Dickie. In this scene, Greenleaf fails to take seriously Marge's (accurate) suspicions

about Ripley's involvement in the disappearance of Dickie. For instance, when Marge discovers a ring that she gave to Dickie in Ripley's room and presents it to Greenleaf as evidence of Ripley's involvement, Greenleaf responds: "Marge there's female intuition, and then there are facts".

Fricker argues that, in responding to Marge in the way that he did, Greenleaf does Marge an injustice; he wrongs her in her capacity as a knower and provider of information. According to Fricker, this undermines her with respect to certain capacities that are essential to her humanity:

> In all such injustices, the victim is wronged in her capacity as a knower. To be wronged in one's capacity as a knower is to be wronged in a capacity essential to human value. When one is wronged or otherwise undermined in a capacity essential to human value, one suffers an intrinsic injustice. The form that this intrinsic injustice takes specifically in cases of testimonial injustice is that the subject is wronged in her capacity as a giver of knowledge.[150]

This form of injustice is clearly a harmful and destructive feature of our social-epistemic practices. As we've seen in the above example, a person can be subjected to this form of injustice as a result of a hearer's prejudicial attitudes toward their gender, but this can also arise out of prejudicial attitudes towards a person accent, social status, race, sexual orientation, etc.

How do these harmful and unjust aspects of our social-epistemic practices bear on the discussion at hand? What I want to argue is that the existence of certain illegitimate *moves* made within a practice doesn't undermine the legitimacy of the *practice as a whole.* Consider the following (simplified) analogy: The fact that there are bad chess moves doesn't somehow call into question the legitimacy of the game of chess itself. In making this analogy, I do not mean to

[150] Fricker (2007) p. 44.

trivialize the harms that are involved in instances of epistemic injustice; the sense in which instances of epistemic injustice are "bad" is clearly not the sense of "bad" that we use when assessing sub-par moves made within a game. My point is just that, even though instances of testimonial injustice represent (in some yet-to-be-fleshed-out sense) *bad* ways of modifying trust vis-à-vis a person, it doesn't follow from this that our practice of epistemic accountability, *taken in its entirety*, is somehow unjust or illegitimate.

This, however, gives rise to the following concern: By what right can I classify instances of epistemic injustice as "bad" moves within our practice? As I have noted, I am trying to answer the justificatory question vis-à-vis epistemic norms in general. My general strategy for answering this question involves attending to our actual social practices to see if there is a distinctive form of response that is taken up in reaction to violations of norms of the relevant kind. My thought is that, if we can answer the justificatory question in this way, that will help us with the content question vis-à-vis epistemic norms (and evidential norms in particular). As I argued in the last chapter, the response captured in **Epistemic Accountability** meets the three criteria for our desired response R. The second condition of that account was as follows: (ii) A modifies her level of trust in B in ways X. However, if I'm simply *starting from* the perspective of the person A who is modifying trust in these ways vis-à-vis B, and, from there, *moving to* an answer to the content question vis-à-vis epistemic norms, then it seems that (given the existence of testimonial injustice) I will be forced to codify norms which license the kinds or reaction that Greenleaf had vis-à-vis Marge. This, however, is clearly unacceptable.

The first thing I would like to note in response to this worry is that, in making this

"move" from individuals who takes up the responses captured in **Epistemic Accountability** to the content of norms, we aren't completely cut off from our wider normative commitments. In other words, this is not a move that is carried out "from nowhere"; it can be guided and constrained by various other commitments. In particular, when it comes to the phenomenon of testimonial injustice, *moral* commitments will come into play. I do not have the space here to adequately explore the kind of harm that is involved in cases of testimonial injustice, but it seems clear that, in taking up the responses captured in **Epistemic Accountability** in the way that Greenleaf does vis-à-vis Marge, he is failing to accord her the level of respect and regard that she is owed. In other words, he is guilty of *wronging* Marge. This is one way, then, of grounding the claim that instances of testimonial injustice constitute "bad" moves within the practice; they are *morally* bad. However, once again, just because there are certain morally bad *moves* that are made within a particular practice, that doesn't mean the *practice as a whole* is illegitimate or unjust; there still might be an abundance of perfectly legitimate moves that are made within the practice. Consider, by way of analogy, medical practice: There might be doctors who are guilty of immoral conduct in their capacity as physicians (e.g. gross negligence, prescribing opioids for pay, etc.). This, however, does not render illegitimate medical practice *as a whole*.

There are two other points that I would like to make in response to this worry. First, recall that I proposed beginning with a few ordinary and commonsense norms as our "starting materials". These included prohibitions on things like believing bald inconsistencies, wishful thinking, and believing on the basis of insufficient evidence. One point I was concerned to make in Chapter 2 is that our starting position vis-à-vis the content of epistemic norms perhaps isn't *so* bleak that we don't have some kind of fuzzy, pre-philosophical handle on their content.

However, I also noted that the above norms could be developed in various theoretical directions, and that they are amenable to different conceptions of the "epistemic". Nevertheless, I think that these starting materials give us at least *some* kind of guidance when it comes to what "epistemic" norms should look like. The kind of norm that would license Greenleaf's behavior vis-à-vis Marge would be so far removed from this starting conception that we could simply dismiss it outright. This, then, is one further way of grounding the claim that instances of testimonial injustice constitute "bad" moves within the practice; they give rise to norms which fail to obey our beginning constraints (however vague or imprecise these might be) on what "epistemic" norms should look like.

Finally, if we return again to the second condition of **Epistemic Accountability** ("A modifies her level of trust in B in ways X"), and specifically to the relevant "ways X", we should note again that such modifications involve ceasing to take a person's "claims" at face value. As I explained in the last chapter, "claiming" that p involves presenting p as true through some kind of expressive act (whether verbal or otherwise), often with the intention of getting others to accept as true the proposition which is presented as true. Thus, we could say that the responses captured in **Epistemic Accountability** represent ways of responding to "claim makers". The class of "epistemic" norms could thus be thought of as delivering various criteria specifying what s*uccessful* "claim makers" look like. Now, as I've noted, if we're simply starting from the perspective of a person A who is modifying her trust vis-à-vis B, and from there, moving to an answer to the content question vis-à-vis epistemic norms, then our success conditions for "claim makers" will be quite skewed, e.g. they might end up counting Marge as an *unsuccessful* claim

maker. Any attempt to block this kind of result might appear ad hoc given my overarching strategy. However, as I've argued above, there are various ways that we can block this move that aren't ad hoc. Moreover, keeping these above points in mind, we can offer one final way to ground the judgment that instances of epistemic injustice represent "bad" moves within our social-epistemic practices.

Consider, once again, **Epistemic Accountability**:

Epistemic Accountability: A holds B epistemically accountable iff:

(i) A judges that B violates a norm of the relevant kind; and

(ii) A modifies her level of trust in B in ways X; and

(iii) A makes the modifications in (ii) because of her judgment in (i).

Say that Greenleaf responds to Marge in the ways captured in condition (ii). In order to count as holding Marge "epistemically accountable", Greenleaf also has to make these modifications because he judges that Marge violates a norm "of the relevant kind". As we've seen, my hope is to fill out these norms by attending to our actual practices. As I've also noted, it seems that norms "of the relevant kind" will specify criteria for successful "claim-makers". The worry we've be considering is that, given my overarching strategy, I'll be forced to codify illegitimate criteria given the reality of epistemic injustice. However, as I've argued above, when it comes to filling out the content of the relevant norms, we aren't hostage to our existing practices: If we end up with norms that license immoral activity, then we can reject those norms; if we end up with norms that deviate too far from our "starting materials", then we can reject those norms, etc.

But now we can discern a third and final way to ground the judgment that instances of epistemic injustice represent "bad" moves within our social-epistemic practices. If we focus once

more on Greenleaf's responses vis-à-vis Marge, and we keep in mind that we can *rule out* norms that might violate the above constraints, then what kind of norm might we be left with that Greenleaf is holding Marge to? It seems that it would have to be one that prohibits believing in ways that are *unreliable* or *intellectually irresponsible*. However, this means that Greenleaf's judgment that Marge violates a norm of the relevant kind will *itself be false*; he is mistaken in this judgment. Marge does *not* believe in ways that are unreliable or intellectually irresponsible, quite the opposite. Moreover, Greenleaf's judgment to the contrary clearly seems to be *culpable.* Not unlike the case of Tom in **Case 1** from the previous chapter (the person who culpably denies the existence of human-induced climate change), Greenleaf is manifesting a bad "doxastic self". In Greenleaf's case, he is both failing to relate to others as he should qua believer, *and* he is failing to respond appropriately to available evidence. It is this latter point – when coupled with the points made above – that can provide us with a third way of grounding the claim that instances of epistemic injustice represent "bad" moves within our social-epistemic practices.[151]

Thus, I do not think the problem of epistemic injustice renders our practice of epistemic accountability illegitimate. Instances of epistemic injustice represent bad *moves* that are made within our practice of epistemic accountability. However, the existence of bad moves within a practice doesn't threaten the legitimacy of the practice *as a whole*. Above I suggested a number of ways to ground this judgment that instances of epistemic injustice represent "bad" moves within our practice that are compatible with the approach to the justificatory question vis-à-vis

[151] Note as well that this means that *others* can subject Greenleaf to the responses captured in **Epistemic Accountability**, perhaps in a way that constitutes "epistemic blame". Greenleaf is also an appropriate target of *moral* accountability responses given his treatment of Marge.

epistemic norms that I've adopted here.

3. Non-Culpable Believers and Epistemic Accountability

The second challenge to the legitimacy of our practice of epistemic accountability concerns non-culpable believers. Recall that, according to the arguments of the preceding chapter, there are certain instances in which we hold non-culpable believers accountable for their doxastic attitudes. I provided the example of Bill in **Case 2**; the non-culpable denier of human-induced climate change. I argued that we can take up the responses captured in **Epistemic Accountability** towards Bill in a way that goes beyond mere assessment vis-à-vis a standard even though Bill is in no way *blameworthy* for his beliefs. One might accept my interpretation of our practices here but then argue that this aspect of our practice is somehow *unjust* or *unfair*. After all, it's not *Bill's* fault that he's in the evidential situation that he finds himself in, and he's doing his best to try to form true beliefs about the subject matter. Given that taking up the responses captured in **Epistemic Accountability** goes beyond mere assessment vis-à-vis a standard (i.e. it is a significant form of response), one might think that taking up these responses to Bill and others like him isn't fair. I'll consider four different attempts to pursue this challenge. I will argue that none of these attempts is successful.

3.1 Accountability and Avoidability

According to this first challenge, subjecting a person to an accountability response is unfair or unjust if the target of the response didn't have a reasonable opportunity to avoid incurring it. Watson (2004) provides a formulation of a general principle which captures this idea in terms of the notion of "sanctions":

It is unfair to impose sanctions upon people unless they have a reasonable opportunity

to avoid incurring them.[152]

While I don't think it's quite right to conceive of epistemic accountability as a "sanction"[153], we can apply this idea to epistemic accountability. Specifically, one might argue that since Bill in **Case 2** didn't have a reasonable opportunity to avoid incurring the responses captured in **Epistemic Accountability**, it is not appropriate to subject him to them.

The thought here is that, since Bill is a blameless victim of his circumstances, it wouldn't be fair to hold him accountable for his beliefs. We should begin by reminding ourselves that the form of response captured in **Epistemic Accountability** is a *non-moral* form of accountability. Consider, for instance, Watson's discussion of the "victim-criminal". This is someone who is a violent criminal but who was also a victim of an abusive childhood. Watson says that, knowing the victim-criminal's background, our tendency to hold him *morally* accountable for his conduct (say, by subjecting him to reactive attitudes such as resentment and indignation) will perhaps be "inhibited."[154] Nevertheless, this is not so when it comes to holding the victim-criminal *legally* accountable:

> To be sure, this inhibition does not shield victim-criminals from legal sanctions. We still protect ourselves against their murderous assaults; we hunt them down, lock them up, shoot them. Hence our scruples about fairness are of no consolation (or compensation) to them. Nevertheless, these concerns affect our sense of what we're doing. Seeing the criminal as himself a victim will not prevent us from shutting the cage or pulling the trigger. But these responses will then tend to seem regulative rather than retributive. In a disconcerting way, they lose their normal expressive function.[155]

There are a number of parallels and differences between Watson's case of the victim-criminal

[152] Watson (2004) p. 276.
[153] I'll return to this issue below.
[154] Smith (2019) disagrees with this and argues that the individuals like the victim-criminal and psychopaths can be appropriately subjected to reactive attitudes.
[155] Watson (2004). p. 281.

and the case of Bill in **Case 2**. As I've mentioned, Bill might also be thought of as a victim of his circumstances. Bill's circumstances are of course not as bad as those endured by Watson's victim-criminal; Bill is "just" being systematically misled when it comes to a certain subject matter. But he is a victim of these circumstances nonetheless. Also, while Bill's circumstances don't lead him to carry out heinous acts like the victim-criminal, he *is* led to take on a host of false beliefs regarding an important subject matter, beliefs that he couldn't reasonably avoid. As a result of all of this, one might argue that it isn't fair to *blame* Bill or to hold him *morally* accountable for his beliefs. Nevertheless, as I have argued, there *is* a form of *non-moral* accountability that we can take up towards Bill, viz. we can take up the responses captured in **Epistemic Accountability.**

A similar point is made by Watson when it comes to holding the victim-criminal *legally* accountable. However, as Watson noted, in the case of the victim-criminal these responses will be "regulative rather than retributive" and that they "lose their normal expressive function." The case of Bill perhaps parallels the case of the victim-criminal insofar as they both show us that there can be *non-moral* forms of accountability that are appropriate even when moral accountability is off the table (and even when the target *didn't* have a reasonable opportunity of avoiding the non-moral forms of accountability). However, there are also important differences between the two cases. Specifically, when it comes to reacting to Bill by taking up the responses captured in **Epistemic Accountability,** it is not the case that these responses will simply be "regulative". Again, this seems to be the view that Boult must take when it comes to the case of Bill. However, as I argued in the last chapter, this seems to get the case badly wrong.

Note again the following three features of the case: (i) Bill's beliefs about climate change do not express or disclose anything objectionable about his "doxastic self"; (ii) Bill is still "answerable" for his beliefs regarding climate change (i.e. he could still be appropriately asked to justify these beliefs); and (iii) Bill's beliefs regarding climate change are *responsive* to new evidence, i.e. were he to acquire strong evidence that his current views were wrong he would revise his beliefs. A purely "regulative" response is perhaps taken up towards someone who is (either temporarily or permanently) incapacitated or otherwise impaired. These are, we might say, "damage control" responses; responses that are taken up for the sake of *management.* As I have argued, this is not the right way to think of our responses vis-à-vis Bill. While the practice of epistemic accountability *as a whole* might perform some kind of regulatory function (e.g. preventing the spread of false information), an *individual's* taking up the responses captured in **Epistemic Accountability** vis-à-vis Bill isn't an entirely "detached" or "regulatory" form of response; we still view Bill as a perfectly rational and reasonable person who is in full possession of his faculties and manifests an exemplary "doxastic self".

Given these features of the case, one might think that the appeal to Watson's victim-criminal actually *doesn't* help when it comes to responding to the present worry. We could reformulate that worry as follows:

> The *only* time it's fair to subject an individual to accountability responses when they lacked a reasonable opportunity to avoid incurring them is when these responses are taken up *solely* for the purpose of regulation and management given that the target of the responses is (either permanently or temporarily) incapacitated or otherwise impaired.

The type of "impairment" at issue might include the situation which has befallen Watson's victim-criminal; this person might be an *incorrigible* wrong-doer given their unfortunate

upbringing. However, as we have seen, Bill is *not* incapacitated or impaired in any way. Thus, it doesn't seem like the responses that we take up in response to him (i.e. those captured in **Epistemic Accountability**) will be based purely in considerations of regulation and management. However, as we've noted, Bill also lacked a reasonable opportunity to avoid incurring these responses. The above reformulated principle implies that taking up these responses towards Bill is unfair. However, the above principle *doesn't* seem to imply this when it comes to holding Watson's victim-criminal legally accountable for his heinous acts. In that case, the individual perhaps *is* "impaired" in the relevant sense, and the legal accountability responses *do* seem to be based purely in considerations of regulation and management. Thus, perhaps appealing to Watson's victim-criminal does not help when it comes to responding to this worry.

The question we need to consider is the following: What is the rationale for accepting the above reformulated principle? And in particular, whatever this rationale might be, does it preserve the verdict that it's inappropriate to subject Bill to the responses captured in **Epistemic Accountability**? I'll consider two possible rationales: the *sanction*-based rationale, and the *communication*-based rationale[156]. I will argue that neither of these approaches preserves the verdict that it's inappropriate to subject Bill to the responses captured in **Epistemic Accountability.**

Let's consider the sanction-based rationale. According to Smith (2019),

[156] I have been influenced here by Smith (2019). Smith is interested in the rationale supporting the judgment that the Strawsonian reactive attitudes (e.g. resentment and indignation) cannot be appropriately taken up towards "moral incorrigibles" (e.g. Watson's victim-criminal) and psychopaths. She considers the two possible rationales mentioned here and finds them wanting. One might also appeal to these rationales in an effort to support the reformulated principle.

> The most common understanding of a sanction is that it is a punitive measure that is deliberately instituted in order to discourage particular forms of behavior.[157]

There are certain kinds of demands that we place on one another that are backed up by sanctions. Being liable to sanctions means that violation of some particular demand or expectation will be met with certain kinds of punitive measures. Sanctioning a person thus clearly goes beyond mere assessment vis-à-vis a standard. Accordingly, considerations of fairness will come into play in instances where a person is sanctioned. Thus, one might think that the justification conditions for sanctions are more stringent than non-sanctioning responses, especially when the individual being sanctioned lacked a reasonable opportunity to avoid the sanction. Specifically, one might flesh out our above principle in the following way:

> The *only* time a *sanction* can be imposed on a person when that person lacked a reasonable opportunity to avoid incurring it is when the sanction is imposed *solely* for the purpose of regulation and management given that the target of the responses is (either permanently or temporarily) incapacitated or otherwise impaired.

This way of fleshing out the above principle would perhaps deliver the verdict that we can impose certain legal sanctions on Watson's victim-criminal. But is this fleshed out principle true? And, importantly, does it deliver the verdict that it's inappropriate to subject Bill to the responses captured in **Epistemic Accountability**? First, there's good reason to think that this fleshed out version of the principle is *not* true. Consider, for instance, cases of "strict liability" in the law. Such cases involve individuals who are liable to sanctions even though they've manifested no ill-will, negligence, or objectionable character traits. Moreover, these cases can involve individuals who are *not* incapacitated or otherwise impaired. Certain traffic violations illustrate this phenomenon. Say that a person is driving with a faulty turn signal or headlight.

[157] Smith (2019). p. 91.

Perhaps the person has no evidence of their vehicle's problem, and they are a very cautious and safe driver (say that they do a thorough check of their vehicle before every drive, but mid-drive some wiring mishap occurs). Vehicle defects like these can escape our attention even if we're highly attentive and conscientious drivers. However, we are still liable to sanctions when such defects occur.[158] Thus, here's a case in which a person is liable to a sanction even though they lacked a reasonable opportunity to avoid incurring it, yet the target of the response *isn't* incapacitated or otherwise impaired.

We thus have good reason to think that the above, sanction-based, way of fleshing out our principle is false. However, even if it *were* true, it wouldn't preserve the verdict that it is inappropriate to subject Bill to the responses captured in **Epistemic Accountability**. This is because it is *not* correct to conceive of these responses as a "sanction". Recall that, on Smith's construal, a sanction is a "punitive measure that is deliberately instituted in order to discourage particular forms of behavior." Recall that one can take up the responses captured in **Epistemic Accountability** vis-à-vis Bill *privately*. In other words, one can make these attitudinal adjustments towards a person without overtly expressing them, whether to the target of the responses or to anyone else. In such cases, these adjustments will clearly not be a way to "discourage particular forms of behavior" on the part of others. If Bill has no idea that others are taking up these responses towards him, then he won't be discouraged from believing what he

[158] There are also cases of product liability. Imagine that a toy manufacturer subjects a certain toy to rigorous safety tests prior to putting it on the market. Say that, even though many steps were taken to ensure the safety of the toy, some issue slipped through the cracks and a child was injured as a result of some undetected design flaw. In a case like this, the toy manufacturer is still liable to sanctions.

does. While there will perhaps be instances where overt expressions of reproach or criticism amount to "sanctioning" a person for believing in a certain way, these overt expressions are not necessary components of the responses captured in **Epistemic Accountability**. Thus, the sanction-based formulation of the principle does not preserve the verdict that it's inappropriate to subject Bill to the responses captured in **Epistemic Accountability** since these responses do not amount to a sanction[159].

Is there another way of fleshing out the above principle that preserves the verdict that Watson's victim-criminal can be appropriately subjected to certain accountability responses but which rules out subjecting Bill to the responses captured in **Epistemic Accountability**? Consider, for instance, the "communication-based" rationale. According to this rationale, given that Watson's victim-criminal lacks certain *communicative capacities*, he is *only* liable to sanctions based in considerations of management and prediction. The thought here is that, other forms of accountability (e.g. moral accountability responses including the Strawsonian reactive attitudes) involve a communicative aspect that is inappropriate to take up towards someone like Watson's victim-criminal. Specifically, one might argue that moral accountability responses like these involve certain *demands* that victim-criminals aren't capable of caring about or comprehending. As Smith (2019) explains:

> It might seem that psychopaths, for example, cannot be appropriate targets of the reactive attitudes, since they are not capable of recognizing or caring about specifically moral reasons. Some have also argued that incorrigible wrongdoers or those living in morally corrupt cultural circumstances may be inappropriate targets of accountability blame, if

[159] This is one important respect in which my view differs from Kauppinen (2018). Kauppinen *does* construe the kinds of trust modification involved in **Epistemic Accountability** as a sanction.

> they are genuinely incapable of recognizing or appreciating the moral concerns communicated by this form of moral response.[160]

The communication-based rationale gives us another possible way of fleshing out the above principle. We could formulate it as follows:

> The *only* time it's fair to subject an individual to accountability responses (whether sanctioning responses or otherwise) when the individual lacked a reasonable opportunity to avoid incurring them is when these responses are taken up *solely* for the purpose of regulation and management given that the target of the responses *lacks certain communicative capacities.*

Without the inclusion of "whether sanctioning responses or otherwise" we would run into the issue discussed above; the responses captured in **Epistemic Accountability** don't amount to a sanctioning response. The relevant kinds of "communicative capacities" will involve the ability to care about or understand the demands that are implicit in the relevant forms of accountability. This principle, if true, would perhaps preserve the verdict that Watson's victim-criminal can be subjected to legal responses. However, it seems that it would *rule out* the verdict that it's appropriate to subject Bill to the responses captured in **Epistemic Accountability**. First, as I've noted, when we take up these responses vis-à-vis Bill, it does not seem that they are rooted in considerations of management and prediction. Second, Bill doesn't lack the relevant communicative capacities. He is fully capable of caring about and understanding the responses captured in **Epistemic Accountability**. As I've mentioned, we can even say that Bill's beliefs are *responsive* to new evidence; *were* he to confront strong evidence which shows that his beliefs are mistaken, he would respond to the evidence by revising his views.

Nevertheless, is the communication-based version of our principle true? Once again,

[160] Smith (2019) p. 101. It's important to note that Smith disagrees with the claim that psychopaths and "moral incorrigibles" cannot be subjected to Strawsonian reactive attitudes.

there is good reason to think that it isn't. First, the communication-based formulation of our principle appears to be ad hoc. Let's say that Watson's victim-criminal doesn't possess the requisite kinds of communicative capacities that would make certain moral demands appropriate. *When it comes to cases like this,* one might hold that accountability responses (whether sanctioning or otherwise) that couldn't reasonably be avoided must be based in considerations of management and prediction. But why think that the *only* time it's fair to subject an individual to an accountability response (whether sanctioning or otherwise) that couldn't reasonably be avoided is when the target lacks certain communicative capacities, and the responses are based in considerations of management and prediction? The above communication-based formulation of our principle seems tailor-made to rule out the appropriateness of subjecting Bill to the responses captured in **Epistemic Accountability**. Why not instead see the case of Bill as a counterexample to the principle? Additionally, the communication-based formulation of the principle confronts the instances of "strict liability" discussed above. These cases involve individuals who are subjected to certain accountability responses that they couldn't reasonably avoid, yet they do *not* lack the relevant communicative capacities.

Thus, I do not think that concerns related to avoidability rule out subjecting Bill to the responses captured in **Epistemic Accountability**. I began by reiterating the point that these responses constitute a *non-moral* form of accountability. I discussed the case of Watson's victim-criminal in order to illustrate the possibility of subjecting a person to a non-moral form of accountability even when the person lacked a reasonable opportunity to avoid incurring it. I also noted that one could argue that there are important disanalogies between Watson's victim-

criminal and Bill in **Case 2**. I went on the consider a number of different ways to formulate principles which could possibly preserve the verdict that it is appropriate to hold Watson's victim-criminal accountable while ruling out the appropriateness of subjecting Bill to the responses captured in **Epistemic Accountability**. I argued that none of these formulations succeed.

3.2 Accountability and Voluntary Control

The second way of pursuing the idea that it would be unfair to subject Bill to the responses captured in **Epistemic Accountability** involves considerations related to voluntary control. One might think that, when it comes to holding a person accountable for their thinking or conduct, the person must exercise some measure of control over the thing that they're being held accountable *for*. When it comes to epistemic accountability, individuals are being held accountable for their doxastic attitudes. Thus, one might hold that, in order for the practice of epistemic accountability to be legitimate, we must exercise some measure of control over our doxastic attitudes. Since it seems rather implausible that we exercise voluntary control over our doxastic attitudes, the practice of epistemic accountability isn't legitimate.

Once again, it's important to remember that **Epistemic Accountability** represents a *non-moral* form of accountability. There are other forms of accountability (e.g. legal) that don't seem to require voluntary control. However, in many cases where legal accountability is called for in the absence of voluntary control, it may be serving a purely "regulatory" function. As I've argued, this isn't the right way to think about the case at hand, viz. the case of Bill in **Case 2.** This might lead one to think that, in the case of Bill and others like it, some kind of control condition must be met in order for the accountability responses to be appropriate.

If there is a control condition which needs to be met in order for the responses captured in **Epistemic Accountability** to be appropriately taken up vis-à-vis Bill, what is that condition? If that condition is *direct voluntary control*, then it will perhaps not be met[161]. It does not seem that we can "believe at will". For instance, I can raise my right arm at will. We seem to exercise direct voluntary control over bodily movements like this. However, it does not seem that I exercise this level of control when it comes to the things that I believe. I cannot, for instance, believe that the moon is made of cheese "just like that". For instance, if someone offered to pay me a million dollars right now to believe that the moon is made of cheese, it doesn't seem that I could thereby take on the belief. Upon being offered the million, I might pursue various courses of action in an attempt to *bring about* or *induce* the relevant belief. But it doesn't seem that I could spontaneously start believing that the moon is made of cheese in the way that I can just spontaneously raise my right arm (although I might *act as though* I believe that the moon is made of cheese in order to get the money).

Thus, if direct voluntary control is a condition for the appropriateness of the responses captured in **Epistemic Accountability**, then those responses will rarely – perhaps *never* – be appropriate. Notice that this will call into the question the appropriateness of those responses *in general*, not just their appropriateness in the case of Bill in **Case 2**. I do not, however, think this represents a serious worry. I see no reason for imposing such a stringent control requirement on

[161] Alston (1988) famously utilized this point to argue "deontological" conceptions of epistemic justification. Alston appealed to a version of the "ought implies can" principle; deontological notions such as "ought" are applicable to something only if it is under our voluntary control. Since beliefs are not under our voluntary control, deontological conceptions of epistemic justification are mistaken.

the form of responses captured in **Epistemic Accountability**, including in the case of Bill. It doesn't seem that we impose such a strict control requirement even when it comes to more "weighty" forms of accountability, e.g. *moral* accountability. For instance, someone can be blamed for some state of affairs even if they lacked the ability to bring about the state of affairs "just like that" through a decision or spontaneous act of will. For instance, I can blame someone for the mess that they made in my apartment while I was out of town. A mess in my apartment can't be brought about "just like that"; there are various intervening steps that have to be carried out in order to bring about the mess. Nevertheless, I can still blame someone for the mess that they made.

Similar points might be made when it comes to our doxastic attitudes. Even if we lack direct voluntary control over our beliefs, there still might be some connection between our choices and decisions and our doxastic attitudes, e.g. perhaps our doxastic attitudes can come about as a result of our *prior* choices, or perhaps we can choose to *endorse* or *identify with* our current doxastic attitudes, or maybe we can bring about *changes* to our current beliefs by choosing to engage in certain actions, etc.[162] Thus, there are various ways in which a control requirement might be met when it comes to our doxastic attitudes.

However, we could even go a step further here and note the possibility of *denying* a control condition altogether when it comes to the appropriateness of accountability responses (whether moral *or* epistemic). For instance, some authors emphasize the possibility of *revealing* or *expressing* ourselves through our thinking and conduct where this is *not* a matter of exercising

[162] For a discussion regarding various "volitional" views of responsibility for our attitudes, and an argument for a "non-volitional" alternative, see: Smith (2005).

control over our thoughts, feelings, and actions. Smith (2005) argues that we express our agency "not only in our explicit choices and decisions, but also in what we unreflectively think, feel, desire, and notice."[163] Smith argues that we often take an individual's spontaneous attitudes, reactions, and unreflective patterns of awareness to be expressive of their agency. These aspects of the person, no less than their voluntary choices and decisions, tell us something about *who the person is* and can appropriately give rise to various kinds of assessments and reactions (including moral responses). As Smith notes, we commonly infer from the presence or absence of these unreflective patterns to claims about what a person *really* cares about or judges to be important. Such attitudes might fail to have *any* connection to our choices or decisions. Nevertheless, they seem to be appropriate objects of moral assessment and criticism.

Thus, I do not think that worries related to voluntary choice and control threaten the appropriateness of the response captured in **Epistemic Accountability**. As I've noted, even if we lack direct voluntary control over our doxastic attitudes, these attitudes can still be connected to our choices and decisions in various ways. Additionally, there are plausible views which deny *any* control condition when it comes to accountability responses (whether moral or epistemic). There are thus various ways to establish that a control condition for the appropriateness of the responses captured in **Epistemic Accountability** – if one exists – is indeed met. There are also ways to deny a control condition altogether.

3.3 Accountability and Blame

One might take the forgoing to open up the possibility of the next challenge. To illustrate

[163] Smith (2005) p. 263. Hieronymi (2008) also a develops an account of "voluntariness" according to which beliefs are not under our voluntary control but we are responsible for them nevertheless.

this next challenge, let's remind ourselves of **Case 1** and **Case 2:**

> **Case 1**: Tom has an abundance of available evidence which is such that, were he to consider it more carefully, he should be rationally compelled to abandon his persistent denial in the existence of human-induced climate change. In spite of this, Tom maintains his belief that human-induced climate change isn't real.
>
> **Case 2**: Bill lives in a community run by a group of individuals who have managed to shut-off access to evidence from the outside world pertaining to the existence of human-induced climate change. Bill is interested in forming true beliefs about this topic, frequently going to the library to research it. Nevertheless, given the evidence available to him, Bill believes that human-induced climate change isn't real.

Say that we've answered the worries presented in the last section regarding voluntary control; there is either some control condition that is met in both of these cases, or there is no control condition at all when it comes to taking up the responses captured in **Epistemic Accountability**. One could proceed to try to argue that the responses captured in **Epistemic Accountability** are not appropriate to take up vis-à-vis Bill in **Case 2** given that he is not *blameworthy*. Thus, according to this challenge, it is inappropriate to subject a person to the responses captured in **Epistemic Accountability** when that person isn't blameworthy for believing what they do.

The thought here is that there might be a *necessary* condition for the appropriateness of the responses captured in **Epistemic Accountability** which is met in both **Case 1** and **Case 2**. Perhaps this necessary condition has something to do with our choices and decisions, e.g. the choices and decisions that we make which have an impact on what we believe. Alternatively, this necessary condition might have *nothing* to do with our choices and decisions, e.g. it might just concern the reasons that we take to support the attitudes that we hold. Either way, according to the current challenge, this is just a *necessary* condition for the appropriateness of the responses captured in **Epistemic Accountability**. In particular, the proponent of the current challenge will

argue that, in the case of Bill in **Case 2**, there is a further condition that needs to be met in order to for these responses to be appropriate, viz. Bill has to be *blameworthy* for believing as he does. Since this condition is met in **Case 1** but not in **Case 2**, it's not appropriate to subject Bill to the responses captured in **Epistemic Accountability**.

There are two points that I would like to make in response to this worry. The first has to do with the concept of "blame". It seems that there will be certain instances where we can clearly take up the responses captured in **Epistemic Accountability** towards a person, yet where it isn't clear whether the concept of "blame" is applicable. Let's return to the case of Tom in **Case 1**. It seems that we can take up the responses captured in **Epistemic Accountability** vis-à-vis Tom in **Case 1**. One might be tempted to follow Boult and call this an instance of "epistemic blame". However, I think part of the temptation here is traceable to the *significance* or *importance* of the subject matter in question (viz. climate change). Consider, by way of contrast, the following case:

> TOM AT THE MOVIES: Tom saw a movie at the downtown cinema this past Tuesday at 7:30 p.m. On the basis of this, Tom believes that the same movie will be playing at the same time at the downtown cinema the following Tuesday. In other words, without checking to see if the movie is even still playing at the downtown cinema the following week, or whether it's still playing but the screening days/times have changed, etc., Tom comes to believe, just on the basis of his seeing the movie on Tuesday at 7:30 p.m., that the same movie will be screened again the following Tuesday at the same time.

Tom lacks sufficient evidence for his belief. Arguably, he is being intellectually irresponsible in maintaining the belief; he is being hasty and careless in believing what he does. As a result, Tom can be seen as an unreliable source when it comes to this particular subject matter, viz. the screening days and times for the particular movie in question at the downtown cinema. It thus

seems that we can take up the responses captured in **Epistemic Accountability** vis-à-vis Tom when it comes to this subject matter. However, does doing this really amount to *blaming* Tom? The relative insignificance of this subject matter makes the notion of "blame" appear less applicable. Even if this is the case, we can still subject Tom to the responses captured in **Epistemic Accountability**.

Perhaps the notion of "blame" can be stretched far enough to be applicable in the case of TOM AT THE MOVIES. Even if this is so, the present challenge still confronts the following issue: There are various forms of accountability that are perfectly legitimate to take up towards a person even when that person isn't blameworthy *in any way*. In requiring blameworthiness as a condition for the appropriateness of the responses captured in **Epistemic Accountability**, the defender of the current challenge seems to be assuming a close connection between *moral* accountability and the forms of responses captured in **Epistemic Accountability**. However, as I have stressed, we should not assume that our practice of epistemic accountability mirrors our accountability practices in the moral domain. Consider, once again, the case of DRUNK DRIVER:

> DRUNK DRIVER: [S]uppose it is necessary for complying with the norm of avoiding reckless driving that your blood alcohol level is below 0.1% by volume. You've had a few drinks that a reliable friend told you had very low alcohol content, and you don't feel drunk, so that you're highly justified in believing that your blood alcohol level is below the required level, and you decide to drive home. But in fact you are over the limit, drive erratically, and are caught by the police. In this case, you may well have an excuse from moral sanctions — in the circumstances, it wouldn't be fair to resent you, say. But you plausibly don't have an excuse from legal sanctions — you should pay your fine. The legal sanction doesn't imply that you're morally bad, or that there's something wrong with your character or your will.[164]

[164] Kauppinen (2018) pp. 8-9.

The person in this example isn't blameworthy for their erratic driving. Nevertheless, they can still be subject to certain forms of legal accountability. My point has been that we can think of Bill in **Case 2** in an analogous fashion. Bill is not blameworthy for believing as he does. Nevertheless, he is still appropriately subjected to the response captured in **Epistemic Accountability**. While these responses do not amount to sanctions imposed by the state, they do go beyond mere assessment vis-à-vis a standard.

3.4 Accountability and Attributability

In the preceding three sections, we've been trying to track down a basis for the claim that taking up the responses captured in **Epistemic Accountability** vis-à-vis Bill in **Case 2** is somehow unjust or unfair. We've considered challenges based in considerations of avoidability, voluntary control, and blameworthiness. In each case, the attempt to support the claim that it's unfair to subject Bill to relevant responses has faltered. The last challenge appeals to the notion of *attributability*. Recall that, "attributability" has to do with whether or not a person's thoughts or behavior somehow express or disclose their true commitments or values. One could appeal to the notion of attributability in order to formulate a final challenge to the legitimacy of our practice of epistemic accountability. Specifically, one could argue that taking up the responses captured in **Epistemic Accountability** towards Bill in **Case 2** is unfair since Bill's false beliefs do not express or disclose anything objectionable about his underlying values or commitments. To put it in language that I introduced in the previous chapter, Bill's false beliefs are not expressive or disclosive of an objectionable "doxastic self". By contrast, Tom's beliefs in **Case 1** *do* seem to express or disclose an objectionable "doxastic self"; Tom is being lazy, dogmatic, and closed-minded in believing as he does. As a result, one might think that it's appropriate to

subject Tom to the responses captured in **Epistemic Accountability**, but not Bill.

We can find an argument resembling this in Boult (2017). Consider the case of DUMMIES, which Boult deploys in his argument:

> DUMMIES: Susan is driving in an obstacle course full of life-like dummies, and her objective is to run all of the dummies over. She has been led to believe (by sources she knows are trustworthy) that it is just a game and there are no real people on the obstacle course. Unbeknownst to Susan, however, a real person has snuck onto the course. In the process of doing her best to succeed in the game, Susan runs the person over.

Here's what Boult says about this case:

> It seems to me that not only is Susan blameless (she does not seem to be the appropriate subject of reactive attitudes or sanctions) for running a real person over, Susan is not attributability-responsible for running a real person over (where the event is to be understood under that description). The fact that Susan runs a real person over is not rationally related to her evaluative commitments (whatever they are). As evidence for this, notice that we would not adjust our assessment of Susan's character, make decisions about the kinds of plans we would make with her, or become less inclined to associate with her on the basis of this event. These are precisely the sorts of moral responses that underwrite attributability.[165]

Boult argues that Susan is not attributability-responsible for running over the person. He then relies on principle (A) to support the claim that, as a result, Susan violated no *norm* in acting as she did:

> (A) There is good but defeasible reason to doubt that a norm N has been violated by an agent S if S is not attributability-responsible for doing what N prohibits.[166]

The case of DUMMIES, along with principle (A), could be used to formulate an objection to the claim that it's appropriate to subject Bill in **Case 2** to the responses captured in **Epistemic Accountability**. Recall that the first condition of **Epistemic Accountability** is the following: "(i)

[165] Boult (2017) p. 341.
[166] Ibid. p. 339.

A judges that B violates a norm of the relevant kind." My claim is that we can respond to Bill in the ways captured in **Epistemic Accountability**. If this is correct, then, given the first condition of **Epistemic Accountability**, Bill will have to be in violation of a norm in believing as he does. However, as I've conceded, Bill's false beliefs about climate change are not *attributable* to him; they are not expressive or disclosive of his true values and commitments. Given Boult's principle (A), this means that Bill will *not* have violated a norm in believing as he does. But this means that we cannot take up the responses captured in **Epistemic Accountability** towards him since the first condition of that account will no longer be satisfied.

Clearly what's doing the heavy lifting here is Boult's principle (A). But why should we accept it? Boult provides two general lines of argument in support of principle (A). First, he claims that (A) does better than other principles when it comes to certain cases. Consider, for instance, his principle (B):

> (B) There is good but defeasible reason to doubt that a norm N has been violated by an agent S if S is blameless for doing what N prohibits.[167]

Boult argues that there are certain cases where norm violation and being open to blame seem to come apart. Boult discusses an example involving a psychopath who cheats an elderly woman out of her life savings. Taking up a blame response to the psychopath seems to imply that the psychopath has certain capacities, e.g. an ability to recognize the interests of others as making a valid claim on him. However, given that the psychopath lacks this capacity, taking up a blame response towards him is perhaps inappropriate. However, we can still think of the psychopath as *callous* or *cold*; his actions and attitudes are still attributable to him insofar as they expressive of

[167] Ibid. p. 333.

"who he is", and we can still evaluate him in a certain way. Thus, according to Boult, this is a case in which (B) fails but (A) does not.

However, one could obviously grant these points without accepting *either* (B) or (A). Is there something more positive that could be said in support of (A)? According to Boult,

> The plausibility of the (A) principle is derived from the plausibility of the basic idea that we make normative demands on the actions and attitudes of persons or agents. We do not make normative demands, for example, on inanimate objects or the mechanical proceedings of nature.[168]

Consider this example:

> AVALANCHE: One day, the weather conditions are just right, and a big avalanche destroys a town in the Rocky Mountains.

Is it plausible to say that the avalanche violated a norm in destroying the town? Boult says that it is not. It is, of course, intelligible to say that the avalanche is *responsible* for the destruction of the town. However, it is not plausible to understand this in terms of *attributability*; the avalanche is merely *causally* responsible for the destruction of the town. Return again to the case of Susan in DUMMIES. Boult's claim is that the relation that the avalanche stands in to the destruction of the town is similar to the relation that Susan stands in to her running over a person. Moreover, according to Boult, "just as we do not think it is appropriate to think of the avalanche as violating a moral norm, it is inappropriate to think of Susan as violating a moral norm by running a real person over on the obstacle course."[169]

Whether or not Boult is right in the case of Susan, I do not think he's right about the case of Bill in **Case 2** and other similar cases. Boult is correct to say that we don't place

[168] Ibid. p. 340.
[169] Ibid. p. 341.

normative demands on inanimate objects and the mechanical proceedings of nature. However, he is *not* correct to say that, in cases where a person isn't attributability-responsible for a certain false belief, we do not place normative demands and expectations on the person. As I have argued, we *do* place certain normative demands and expectations on such a person; this is revealed in the fact that we modify our level of trust in them in the ways captured in **Epistemic Accountability**. Moreover, it is simply not true that Bill stands to his false beliefs about climate change in the way that the avalanche stands to the destruction of the town. It is true that Bill's false beliefs aren't *attributable* to him. However, he is still *answerable* for his beliefs, and he is also *responsive* to new evidence[170]. Avalanches and other such phenomena are not answerable for *anything*, nor are they "responsive" in the relevant sense. The above argument is thus unsuccessful. Boult has failed to establish principle (A) as correct, at least when it comes to epistemic matters.

4. Conclusion

In the previous chapter I argued that we have a practice of epistemic accountability. As I've noted, these results could help us answer the justificatory question vis-à-vis epistemic norms in general and, in turn, the content question vis-à-vis evidential norms (including EN). However, an adequate answer to the justificatory question can only be delivered if the social practice of epistemic accountability is a *legitimate* social practice. In the present chapter, I have argued that

[170] We could make similar points regarding Susan. Consider, for instance, the ways in which we might *anticipate* an oncoming avalanche: we might take certain measure to stop it, or to somehow prepare for its unavoidable consequences. This is *not* how we would deal with Susan; we would (urgently) *bring it to her attention* that there's a person in the obstacle course. Only if it was *too late* to engage with Susan in this way would we then take up preventative measures akin to those that we take up in anticipation of the avalanche. If these preventative measures fail and Susan ends up running the person over, we might then expect certain things from Susan (e.g. an apology without an admission of guilt or culpability).

our practice of epistemic accountability is indeed a legitimate social practice. I began by making explicit a certain methodological assumption when it comes to practices and their legitimacy. Specifically, I assumed the following: A social practice X is legitimate if and only if there is no successful challenge to its legitimacy. Importantly, a "successful challenge" doesn't have to be expressed or articulated by any person or group of people; what matters is the *existence* of a successful challenge to the practice. I then went on to consider two general strategies to try to undermine the legitimacy of our practice of epistemic accountability. The first strategy appeals to instances involving epistemic injustice, especially testimonial injustice. I argued that, while instances of epistemic injustice represent *bad moves* made within our social practices, the existence of such moves doesn't undermine the legitimacy of the practice *as a whole*. I then developed a number of ways to ground the judgment that instances of epistemic injustice represent bad ways of modifying trust that are compatible with my overarching approach. After this, I considered a number of attempts to establish that our practice of epistemic accountability is *unfair*. I considered challenges based in considerations avoidability, voluntary control, blameworthiness, and attributability. I argued that all such attempts fail to undermine the legitimacy of our practice of epistemic accountability.

I do not claim that the challenges that I have considered here are exhaustive. However, I do think that the challenges that I've considered represent some of the most forceful challenges to the legitimacy of our practice of epistemic accountability. Given that none of these challenges succeed, I think we have a very good case for the legitimacy of our practice of epistemic accountability. We are now in a position to develop answers to the justificatory question vis-à-

vis epistemic norms in general, as well as the content question. In the next chapter I will turn to these tasks.

Chapter 6: Answering the Justificatory and Content Questions

1. Introduction

Allow me to recap. I began this dissertation by arguing against a view that I labeled "evidential minimalism". According to evidential minimalists, strong evidence E for the truth of p is either *itself* a normative reason for S to believe p, or such evidence *gives rise* to such a reason when certain other "minimal" conditions are met (e.g. S possesses the evidence and considers the question "whether p"). Even though the minimalist view is false, there still might be correct norms of belief that incorporate the notion of evidence. For instance:

> (a) Don't believe without adequate evidence.
>
> (b) Don't believe in a way that flies in the face of the total evidence in your possession.
>
> (c) Don't reach a conclusion without having gone to appropriate lengths to seek out countervailing evidence.

If evidential minimalism is false, then it seems that we need some other way to understand evidential norms of belief. Earlier I set aside the "aim of belief" approach which would have us understand evidential normativity in terms of the attitude of belief itself. I argued that such an approach is unable to successfully account for the fact that evidence sometimes constitutes a reason for us to *expand* our current view. In order to better capture the normativity of *evidence,* I offered the following norm EN:

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.[171]

[171] The relevant sense of "possession" at play in EN is the one that I labeled as "robust" evidence possession in Chapter 1. A subject S possesses strong evidence E for the truth of p in the "robust" sense at time t iff (i) S has, at time t, some contentful mental state (whether occurrent or non-occurrent) which represents E as true, and (ii) S is aware of E's evidential import vis-à-vis p

The question, then, is whether or not EN is a correct norm of belief. I argued that, as an evidential norm, EN is a part of a wider class of "epistemic" norms. Thus, the question is whether or not EN is a correct epistemic norm. I formulated the following "Application-Acceptance Connection" (AAC):

AAC: If a norm N applies to S, then there is good reason for S to accept N.

When a norm N to "applies" to an individual S, that means that it is appropriate for other people to *hold S accountable* for complying or failing to comply with N. When it comes to the norms of *morality* (e.g. norms that require us to keep our promises to others), it seems that such norms have a very *wide* scope of application. I noted that, similarly, it seems that *epistemic* norms (including (a)-(c) above, and perhaps EN) have a fairly wide scope of application. Given AAC, this means that, if N is an *epistemic* norm, then a fairly wide class of individual will have good reason to accept N.

I also distinguished two questions that we can ask about epistemic norms: the "content" question and the "justificatory" question. The content question asks what the relevant norms *are* or what they *say*. The justificatory question asks about our reasons for accepting some norm or class of norms. Thus, when it comes to establishing EN as a correct *epistemic* norm, we need to somehow secure a fairly wide scope of application for the norm. Once again, according to AAC, this means that a fairly wide range of individuals must have good reason to accept EN. In other

---

at t. "Awareness" of E's evidential import vis-à-vis p is an *ability* or *capacity* had by S. Thus, S can be "aware" of E's evidential import vis-à-vis p at time t without explicitly attending to E at t, and without believing p at t. Recall as well that, in order for some evidence E to be "strong" evidence for the truth of p, it must meet the following constraints: E is a true proposition (possibly about the external world), and E makes the truth of p sufficiently likely; it is possible for one to come to *know* p on the basis of E.

words, the reason(s) appealed to in an answer to the justificatory question vis-à-vis EN must be reasons that are shared by wide class of individuals.

In order to establish EN as a correct epistemic norm, I proposed beginning with the justificatory question vis-a-vis epistemic norms in general. My hope is that, if we can provide an answer to the justificatory question vis-à-vis epistemic norms in general, that will help us with the justificatory and content questions vis-à-vis evidential norms (including EN) given that such norms are a part of the wider class of epistemic norms. However, I also noted a significant roadblock confronting my proposed approach: there doesn't seem to be any non-arbitrary way to demarcate the class of "epistemic" norms. The way I proposed overcoming this roadblock is by attending to our actual social-epistemic practices in order to discern a *form of response* that is taken up in reaction to norm violations of the relevant kind. Before pursuing this approach, I argued against a number of alternative ways of understanding "epistemic" norms, including Friedman's "zetetic" approach, and epistemic instrumentalism. I then argued that *there is* a form of response that is taken up in reaction to norms of the relevant kind. I captured these responses in **Epistemic Accountability**. I went on to argue that our practice of epistemic accountability is a legitimate social practice.

We are now in a position to develop an answer to the justificatory question vis-à-vis epistemic norms in general. In the next section I will develop an answer to this question.

2. Answering the Justificatory Question vis-à-vis Epistemic Norms in General

In answering the justificatory question vis-à-vis epistemic norms in general, I will sketch a chain of reasoning that any participant in the social practice of epistemic accountability could engage in; one which reveals the reasons which ground our acceptance of epistemic norms. The

rough idea is as follows: Our social practice of epistemic accountability is structured by various norms. As we have seen, the first condition of **Epistemic Accountability** is the following: "(i) A judges that B violates a norm of the relevant kind". Thus, any time a person subjects another to the responses captured in **Epistemic Accountability**, these modifications in trust will be underwritten by a judgment that the target of the response violates a norm of the relevant kind. My proposal is that *these* are the norms that could be considered the "epistemic" norms. I have also argued that our practice of epistemic accountability is a legitimate social practice. Given this, I will argue that anyone who is a *participant* in the practice will have good reason to accept the norms that structure its activities.

Recall that, when it comes to "epistemic" norms of belief, there was a prima facie presumption that they have a fairly wide scope of application. Given AAC, this means that there will be a wide range of individuals who have good reason to accept such norms. Importantly, AAC does not say that a norm N applies to S only if S *actually accepts* N. Rather, it says that a norm N applies to S only if *there's good reason* for S to accept N. Thus, what we need is a chain of reasoning that a participant in the practice *could* engage in; one which makes explicit what the reasons are which ground one's acceptance of the relevant norms. Recall that this is what the "justificatory" question asks for: the reason(s) which ground an individual's acceptance of some norm or class of norm. "Accepting" a norm or class of norms is an attitudinal state which one takes up vis-à-vis the relevant norm or class of norms. As a result, we can make our acceptances explicit, and articulate the reasons which support them.

I will operate with the following "basing" requirement on normative reasons for Φ-ing: R is a normative reason for S to Φ only if it's possible for S to Φ on the basis of R. The relevant act

Φ for my purposes is accepting a norm or class of norms. Thus, for individuals who haven't *already* accepted epistemic norms, what needs to be shown is that there is some chain of reasoning (i.e. a basis) which could underwrite their acceptance of epistemic norms. For instance, my character of Tom (in **Case 1** and in TOM AT THE MOVIES) has perhaps not accepted epistemic norms. Nevertheless, we can still subject him to the responses captured in **Epistemic Accountability**, i.e. we can hold him epistemically accountable. Given AAC, this means that there must be good reason for Tom to accept the norms implicit in the responses captured in **Epistemic Accountability**. The chain of reasoning below will reveal what those reasons are, and how Tom might reason his way towards accepting the relevant norms on their basis. For individuals who have *already* accepted epistemic norms, their acceptances may be justified even if they haven't engaged in an explicit process of reasoning which reveals the reasons that underwrite their acceptances. It's enough that a process of reasoning like this *exists* and that these individuals *could* engage in it. Note that "could" is different from "*can*"; there might be various counterfactual elements that have to obtain in order for a person to actually be in a position to carry out the requisite process of reasoning. However, these counterfactual elements don't have to do with the actual truth or falsity of the various steps in the chain of reasoning offered below; they have to do with is an individual's ability to *apprehend* the truth of these steps, and their ability to *put them together* in the right way.

When it comes to accepting "epistemic" norms, a participant in the practice of epistemic accountability could reason as follows:

> 1. I hold other people accountable by taking up the responses captured in **Epistemic Accountability**.

2. Holding others accountable in this way is part of an existing social practice.

3. I am thus a participant in the social practice of epistemic accountability.

4. The social practice of epistemic accountability is a legitimate social practice.

5. Participants in a legitimate social practice are bound by the norms and expectations that structure the practice.

6. Therefore, I am bound by the norms and expectations that structure the social practice of epistemic accountability (i.e. such norms apply to me).

7. If a norm N applies to S, then S has good reason to accept N (AAC).

8. Therefore, I have good reason to accept the norms that structure the social practice of epistemic accountability.

9. Therefore, I should accept the norms that structure the social practice of epistemic accountability.

Below I will break down this chain of reasoning step-by-step. However, before doing that, I want to note one thing: The above isn't a mere "existence" proof; it doesn't just establish that *there is* good reason to accept the norms that structure our practice of epistemic accountability while remaining silent on what those reasons *are*. Rather, this answer to the justificatory question *tells us what the reason are* which ground our acceptance of the relevant norms. What *are* these reasons, according to this answer to the justificatory question? Nothing other than the following: these norms structure a legitimate social practice that we participate in. In other words, according to my answer to the justificatory question vis-à-vis epistemic norms in general, it is *not* the case that the relevant norms bind us in virtue of the fact that there's some independent reason that grounds our acceptance of them. Rather, according to my view. we have reason to accept the norms *in virtue of the fact that they are binding on us.* The fact that they are binding on us *is* the reason to accept them. But why are they binding on us? Once again, they are binding on us given

that they structure a legitimate social practice that we participate in[172].

On my view, then, social practices and our status as participants within them are the most basic elements. Social practices give rise to various norms that we can discern, articulate, reject, revise, accept, etc. Our "participation" in such practices consists in the fact that we *hold others* (and perhaps *ourselves*) to the relevant norms. "Holding" each other to norms involves taking up the various responses that are called for when the norms are violated. This will involve different things depending on the type of norm in question. As I have argued, when it comes to "epistemic" norms, this will involve the responses captured in **Epistemic Accountability**. From such practices of interpersonal accountability, we can discern certain norms.

As I have noted, sometimes our practices are *illegitimate* or *unjust* and the norms which structure their activities should be revised or rejected. However, when a social practice is legitimate, there will be no reason to reject or revise the norms that structure its activities. If a person participates in a social practice of this kind, then they will be bound by the norms that structure its activities. I thus do not see any *further* reason that needs to be identified – one which would ground a person's acceptance of some norm or class of norms – once it has *already* been determined that these norms structure a legitimate social practice that the person participates in; the justificatory question has already been answered at that point. As I have made clear, when it comes to the legitimacy of our practice of epistemic accountability, there will be various *constraints* which will come into play when determining the norms that structure its activities.

[172] But isn't "participation" doing a lot of work in explaining why the norms are binding on us? If "participation" is a matter of decision or choice, then perhaps what's really explaining why the norms are binding on us are voluntarist considerations related to decision, choice, or consent. I'll return to this issue momentarily.

However, these constraints do not *confer* legitimacy upon the practice as a whole, they just set *limits* on what a practice's norms can look like. Legitimacy is thus not "conferred" upon our social practice of epistemic accountability by something else; its legitimacy consists in nothing more than the fact that there's no challenge which *succeeds in undermining* its legitimacy.

Keeping these points in mind, let's work through each step of the above answer to the justificatory question. Step 1 says "I hold other people accountable by taking up the responses captured in **Epistemic Accountability**." Step 1 will hold of anyone who modifies their level of trust in others in the ways captured in **Epistemic Accountability**. I take it that a very wide class of individuals takes up the form of response captured in **Epistemic Accountability**. As I noted in Chapter 4, our social-epistemic predicament is such that we invariably rely on each other when it comes to the acquisition and transmission of information. As a result, we also modify our level of trust and reliance in each other when it comes to acquiring information. The responses captured in **Epistemic Accountability** thus represent a core feature of our social-epistemic practices. Specifically, the responses in **Epistemic Accountability** are taken up towards another when that person is deemed to be unreliable when it comes to a certain topic or subject matter. Once again, I take this to be a common feature of our social-epistemic practices; a form of response that many people will engage in.

Step 2 says "Holding others accountable in this way is part of an existing social practice". This step simply makes explicit the fact that modifying trust in the ways captured in **Epistemic Accountability** is part of an existing social practice. Step 3 says "I am thus a participant in the social practice of epistemic accountability". This step says that, insofar as one modifies one's level of trust and reliance in other people in the ways captured in **Epistemic Accountability**, one

is thereby participating in the social practice of epistemic accountability. This seems exactly right. Social practices involve various ways in which individuals interact with each other. The practice that I've focused on here is our social practice of epistemic accountability. It is enough to count as a "participant" in this practice if one actually engages in the relevant form of interaction. Thus, modifying one's trust and reliance in others in the ways captured in **Epistemic Accountability** is enough to make one a participant in the practice. Step 4 says "The social practice of epistemic accountability is a legitimate social practice". This step was established in Chapter 5. I considered a number of forceful challenges to the legitimacy of our practice of epistemic accountability. I argued that none of these challenges succeeded in undermining the legitimacy of the practice. If there's no successful changeful to the legitimacy of the practice, then the practice is legitimate.

Step 5 will require a longer discussion. Step 5 says "Participants in a legitimate social practice are bound by the norms and expectations that structure the practice." The thought here is that, if individuals *partake* in a certain kind of social practice, *and* that social practice is legitimate, then the norms which structure the practice "bind" these individuals. What does it mean for a participant S to be "bound" by a norm N? What this means is that S can be appropriately *held accountable* for complying or failing to comply with N. In other words, other people (or S herself) can appropriately respond to S's violations of the relevant norms by taking up a host of reactions associated with the practice and the norms which govern its activities (where such reactions can include things like sanctions, blame, censure, guilt, reproach, etc.).

It will be important to note a number of features which differentiate our social practice of

epistemic accountability from other legitimate social practices. As a result of these differences, the application of step 5 ("Participants in a legitimate social practice are bound by the norms and expectations that structure the practice") will have a certain significance when the social practice in question is our practice of epistemic accountability. Consider, by way of comparison, club rules, traffic laws, and procedural guidelines of various kinds. These examples make clear the possibility of social practices which incorporate *different* sets of norms, yet which are also equally legitimate. For instance, two different clubs might adopt different sets of norms (one might meet on Wednesday evenings, the other on Thursday evenings, etc.), two different countries might adopt different traffic regulations (one might require driving on the right side of the road, the other on the left, etc.), and two different governing bodies might adopt different procedural guidelines (one might adopt Rules of Order X, the other might adopt Rules of Order Y, etc.). There will thus be an abundance of (relatively "local") practices, many of which will be perfectly legitimate and will also incorporate norms that *differ* from the norms which govern other, equally legitimate, practices of a similar kind.

Is our practice of epistemic accountability like this? In other words, are there different, equally legitimate, practices which are similar in kind to our practice of epistemic accountability yet which operate with *different* norms? Answering this question requires saying something about the answer to the "content" question, which I will turn to below. We can, however, make the following preliminary observation: Given the preceding discussion, it seems that our practice of epistemic accountability will be structured by various *externalist* norms of belief, e.g. norms

that require our beliefs to be *true* or amount to *knowledge*.[173] Recall, once again, the case of Bill in **Case 2**. I've argued that we can appropriately take up the responses captured in **Epistemic Accountability** vis-à-vis Bill. However, Bill is in no way *blameworthy* for believing what he does; in taking up the responses captured in **Epistemic Accountability** towards Bill, we aren't blaming him for the way that he conducts his doxastic life. It seems, rather, that our responses are rooted in the fact that Bill's beliefs concerning climate change are *false*; they don't accurately reflect how things *actually are*. Again, I will spend more time on the "content" question below. However, we can take this as a preliminary case for the claim that our practice of epistemic accountability is structured by various externalist norms of belief, i.e. norms which somehow require our beliefs to accurately reflect the way that things actually are.

Thus, our question is the following: Are there other, equally legitimate, social practices which are similar in kind to our practice of epistemic accountability, but which operate with *different* norms? While there might be some slight variation between different practices of this kind, it seems that, given certain realities pertaining to human needs, interests, and limitations, *any* such practice will incorporate various externalist elements. Consider, for instance, the following observations made by Edward Craig in his *Knowledge and the State of Nature*:

> Human beings need true beliefs about their environment, beliefs that can serve to guide their actions to a successful outcome. That being so, they need sources of information that will lead them to believe truths. They have 'on-board' sources, eyes and ears, powers of reasoning, which give them a primary stock of beliefs. It will be highly advantageous to them if they can also tap the primary stocks of their fellows – the tiger that Fred can see and I can't may be after me and not Fred – that is to say, if they act as informants for each other. On any issue, some informants will be better than others, more likely to supply a true belief. (Fred, who is up a tree, is more likely to tell me the truth as to the

[173] This will require refinement, but it suffices to note at this point that the norms which structure our practice of epistemic accountability will incorporate externalist elements.

> whereabouts of the tiger than Mabel, who is in the cave.) So any community may be presumed to have an interest in evaluating sources of information; and in connection with that interest certain concepts will be in use.[174]

Craig's hypothesis is that the concept of 'knowledge' is one such concept which is bound up with our interest in evaluating sources of information. However, we could also make similar points about the form of response captured in **Epistemic Accountability** and the norms which underwrite those responses. It's important to note that Craig is interested in deriving application conditions for the concept of 'knowledge' by appealing to a hypothetical "state of nature" scenario. The hypothetical "state of nature" scenario is not playing any role in my answer to the justificatory question. Moreover, the instrumental utility of our practice of epistemic accountability and its associated norms is also playing no role in my answer to the justificatory question (although I do not deny that the practice has instrumental utility). Craig's observations simply strike me as plausible, just as contingent facts of the matter regarding creatures like us and the environments that we find ourselves in: Any community of individuals will have an interest in acquiring accurate information about their surroundings. Given certain realities pertaining to human needs, interests, and limitations, any social practice akin to our practice of epistemic accountability will likely incorporate various externalist elements.

Given this, it seems that our practice of epistemic accountability will be different from other legitimate social practices in important respects. First, as I've noted, the possibility of other, equally legitimate, social practices of the same kind which incorporate norms which differ significantly from the ones at play in our practice of epistemic accountability seems rather unlikely given Craig's observations. For creatures like us, in a world like this, it seems that any

[174] Craig (1990) p.2.

practice akin to our practice of epistemic accountability will incorporate various externalist elements. This is thus one respect in which the norms which structure our social practice of epistemic accountability will differ from club rules, traffic laws, and procedural guidelines of various kind. However, there is another important difference that we should also highlight: Club rules, traffic laws, and procedural guidelines are all *enacted*. In other words, these norms are explicitly *chosen and adopted* by groups of individuals who have an interest in coordinating their activities in a certain way. The norms which structure our practice of epistemic accountability are *not* like this. I have argued that those norms are implicit in the responses that we take up towards one another, viz. the responses captured in **Epistemic Accountability**. The norms which structure our practice of epistemic accountability are thus implicit in a certain kind of attitudinal response which is a common feature of our social-epistemic lives. We did not *choose* to be governed by these norms, nor did we *enact* them in some way.

But does this mean that we have *no* choice when it comes to our being bound by epistemic norms? As I've indicated, epistemic norms bind us in virtue of the fact that we participate in a legitimate social practice. Perhaps "participation" is where we can exercise choice when it comes to the binding authority of epistemic norms. After all, it's not the legitimacy of the practice *alone* which accounts for the fact epistemic norms are binding on me, it's also the fact that I'm a *participant* in the practice. Here we need to remind ourselves of what "participation" amounts to when it comes to our practice of epistemic accountability. All it takes to participate in the practice is to take up the responses captured in **Epistemic Accountability** towards another. Can we choose to *not* take up these responses? Yes and no. "Yes" if the

question is whether or not we can, *in particular instances*, decide to temper our responses in various ways, or perhaps not take them up at all, even after correctly judging that a person has violated a norm of the relevant kind (whether culpably or non-culpably). This is indeed a possibility; we are not *incapable* of tempering the responses captured in **Epistemic Accountability**, or abstaining from them altogether, in *individual cases* where we correctly judge that another has violated a norm of the relevant kind.

However, if the question is whether we can choose to *abstain from these responses altogether*, then I think the answer is "no". It's important to remind ourselves here of the fact that the responses captured in **Epistemic Accountability** have to do with modifying trust in a person's "claims", where "claiming" that p involves presenting p as true through some kind of expressive act (whether verbal or otherwise), often with the intention of getting others to accept as true the proposition which is presented as true. "Claims" can be made in many different ways: through explicit speech acts in the course of a face-to-face conversation, over the phone, through text messages, in books, on social media, in letters, through gestures or motions, on post-it notes, etc. The responses captured in **Epistemic Accountability** involve modifying trust in a person's "claims", so understood. This can involve ceasing to take a person's expressed claims that p as good reason to think that p is true. This can also involve a reluctance to seek the person out as a testimonial source of information when it comes to a certain topic or subject (where "testimony" involves making claims). I do not think that abstaining from these forms of response *in their entirety* is a realistic option. I'm doubtful that this is even psychologically possible. For human beings who are engaged in social life with others, responses such as these are an inevitable fact of life. What this means is that "participation" in the social practice of epistemic accountability

will be something which is exceedingly difficult – perhaps even practically *impossible* – to evade. This is thus another important respect in which our practice of epistemic accountability differs from other legitimate social practices; "participation" is practically unavoidable.

We're now in a position to see how the application of step 5 in the above chain of reasoning ("Participants in a legitimate social practice are bound by the norms and expectations that structure the practice") is quite unique when the particular social practice in question is our practice of epistemic accountability. Specifically, our social practice of epistemic accountability seems to be unlike other legitimate social practices in at least the following two respects: There aren't alternative, equally legitimate, practices of the same kind that incorporate significantly different norms, and participation in the practice is practically unavoidable. Given these features, along with steps 1-4, we can see how step 5 generates a rather strong kind of "bindingness" when the social practice in question is our practice of epistemic accountability. Not only will the range of people "bound" by the relevant norms be very wide ranging, it will also be the case that their participation in the relevant practice is practically unavoidable, and the practice itself will be very fundamental to human social life as we know it. Our practice of epistemic accountability could even be considered a basic or foundational human practice; one which makes possible the successful functioning of various other human practices. Given all of this, the norms will "bind" the participants in the practice in a rather strong sense.

Step 6 says "Therefore, I am bound by the norms and expectations that structure the social practice of epistemic accountability (i.e. such norms apply to me)". This is the first-personal recognition that the norms that structure the social practice of epistemic accountability

apply in one's own case. In other words, this step is an acknowledgement that the norms apply to one in virtue of the preceding steps of the argument. To say that the norms *apply* to me is to say that I can be justifiably *held* to the norms which structure the practice (whether by others or myself). As we've seen, when it comes to our practice of epistemic accountability, the way in which we *hold* individuals to the relevant standards is by taking up the responses captured in **Epistemic Accountability**[175]. Thus, step 6 is the first personal recognition that I can be appropriately subjected to these responses (whether by others or myself)[176]. Step 7 is the principle AAC: "If a norm N applies to S, then S has good reason to accept N". As we've seen, this principle states a plausible connection between a norm's "application" to an individual S, and S's acceptance of the norm. The idea is that, if we're *holding* S to a certain norm or standard, then that seems to imply that we think that there's good reason for S to *accept* the relevant norm or standard. "Acceptance" of a norm N by a person S is some kind of attitudinal state that S takes up vis-à-vis N; one which doesn't guarantee S's *conformity* with N (try as we might, external circumstances aren't always congenial) and which is also not gotten simply *via* conforming with N (one can conform to a norm as a result of accident, luck, or coincidence). As I've mentioned, our *holding* a person to some norm or standard seems to imply that we think there's good reason for that person to *accept* the relevant norm or standard. We will be justified in *holding* a person

[175] There are instances in which we "hold" a person to the relevant norms even though we don't actually modify our level of trust in the person in the ways captured in **Epistemic Accountability**. For instance, there might be cases where we *excuse* an individual for their norm violation and abstain from modifying our level of trust in them.

[176] How can I hold *myself* epistemically accountable? Imagine that I came to find out that I'm susceptible to certain fallacious lines of reasoning, or that I'm prone to wishful thinking when it comes to a certain topic or subject matter. After finding this out, it seems that I can modify trust and reliance in *myself*. For instance, I might double check the evidence in appropriate circumstances, or I might cease to "trust my own judgment" in certain cases.

to some norm or standard only if the norm or standard *actually* applies or "binds" the individual. However, when it comes to the norms which structure our social practice of epistemic accountability, we have already established that they do; this was established in steps 1-6. Specifically, the norms that structure our practice of epistemic accountability are binding on us in virtue of the fact that we participate in a legitimate social practice.

Step 8 say "Therefore, I have good reason to accept the norms that structure the social practice of epistemic accountability." This step clearly follows from the preceding steps. Anyone who participates in the social practice of epistemic accountability can come to the first-personal recognition that the norms which structure the practice are binding on them. The first-personal recognition was captured in step 6. Combine this with the principle AAC in step 7 and you get the first-personal recognition in step 8 that I have good reason to accept the norms that structure the practice of epistemic accountability. Once again, this is not a mere existence claim; the chain of reasoning reveals what these reasons *are*. The reason which underwrites our acceptance of the relevant norms is the fact that those norms are binding on us, and they are binding on us in virtue of the fact that we participate in a legitimate social practice. Step 9 says "Therefore, I should accept the norms that structure the social practice of epistemic accountability." On the basis of steps 1-8, one can arrive at the judgment that one should accept the norms that structure our practice of epistemic accountability. Thus, we have a chain of reasoning that anyone who is a participant in our practice of epistemic accountability can engage in. This chain of reasoning could underwrite an individual's acceptance of the relevant norms, even if they haven't *already* accepted the relevant norms.

3. Answering the Content Question

We are now in a position to approach the content question vis-à-vis epistemic norms in general, and evidential norms in particular. Recall that the first condition of **Epistemic Accountability** is the following: "(i) A judges that B violates a norm of the relevant kind". Given the form of responses in the second condition of **Epistemic Accountability** ("(ii) A modifies her level of trust in B in ways X"), we can fill out the content of the norms in the first condition by looking to the circumstance in which we modify our level of trust in the ways captured in the second condition.

First, we should return to the norm EN:

> EN: When you possess strong evidence E for the truth of p at t, and the evidential connection between E and p is clear to you at t, then, at the very least, you have a warranting reason to believe p at t.

As I argued in Chapter 3, epistemic instrumentalists are unable to accommodate this norm. Specifically, all of the versions of instrumentalism that I considered failed to secure a wide scope of application for EN; the appealed to aims, interests, and goals could not secure acceptance of EN for a sufficiently wide range of individuals. As I noted in Chapter 3, an instrumentalist could accept these results and simply argue that EN *isn't* an epistemic norm (or perhaps they could accept that it is, but hold that it only "applies" to a limited range of individuals). In other words, an instrumentalist could argue that, in instances where we *can't* secure an individual's acceptance of EN by somehow appealing to agential aims, interests, and goals, the individual in question *doesn't* have reason to accept EN. Given AAC, it follows that, in such cases, it's not appropriate to *hold this person accountable* for complying with EN.

We are now in a position to see that this position is untenable. Imagine a person who

doesn't accept EN. While this person might treat strong evidence E for the truth of p as a reason to believe p in many instances, they don't accept EN. In particular, they tend to view evidence that doesn't bear on their own personal interests as irrelevant or insignificant. This person is importantly different from the character of Katlyn from Chapter 1. Recall that Katlyn was the individual who refused to *take on* beliefs in pointless and trivial disjunctive claims even though she possessed decisive evidence in support of them (e.g. "Either her name is Katlyn or aliens exist", "Either her name is Katlyn or aliens exist or the moon is made of cheese", etc.). Even if Katlyn doesn't *take on* beliefs in these claim, she can still acknowledge *that there's (warranting) reason* to believe them. I argued in Chapter 1 that Katlyn's isn't appropriately held accountable for not taking on beliefs in pointless disjunctive claims (or other "trivial" truths). However, the person that I'm imagining now doesn't just *fail to take on* certain beliefs that are well supported by her evidence yet which are of no interest to her, *she doesn't even acknowledge that there's good reason to do so*; she views such evidence as entirely irrelevant to the question of what constitutes an adequate basis for belief.

My claim is that we can respond to a person like this by subjecting them to the responses captured in **Epistemic Accountability**. In this particular case, taking up these responses may not involve ceasing to take the person's expressed "claims" at face value. Perhaps this person has various beliefs that are well-supported by good evidence, e.g. beliefs pertaining to issues that are of interest to her. Thus, this person's expressed claims might be reliable and trustworthy. However, say that there are a whole host of topics and issues that this person simply doesn't care about. When this person possesses strong evidence E for the truth of a claim pertaining to such topics and issues, she views it as something which is altogether irrelevant when it comes to the

question of what constitutes an adequate basis for belief. As a result of this, we can subject this person to certain forms of response that are included in **Epistemic Accountability**, e.g. we can manifest a reluctance or unwillingness to seek this person out as a testimonial source of information when it comes to various topics and issues. Moreover, we can *intend* to lower trust in this person's expressed claims. For instance, if she *were* to express some claim regarding a topic or issue outside of her areas of personal interest, we *wouldn't* treat it as good reason to think that the claim is true. Responses such as these are included in the forms of responses in **Epistemic Accountability**.

The above reveals that EN is a part of the class of norms that structure our practice of epistemic accountability. Given the answer to the justificatory question offered above, this means that individuals who are participants in the practice will have good reason to accept EN. Thus, we can now see the full force of the earlier argument presented against epistemic instrumentalists. If my arguments there were on the right track, then instrumentalists cannot secure a wide scope of application for a norm like EN. This means that they must say that, in instances where an individual's acceptance of EN cannot somehow be secured by appealing to agential aims, interests, and goals, the individual has no reason to accept the norm. Given AAC, this means that EN will not "apply" to such individuals. As I have argued, this is deeply mistaken: EN applies to individuals even in instances where the person's acceptance of the norm cannot be underwritten by appealing to agential aims, interests, and goals. Given AAC, there must be a reason for such individuals to accept EN. My answer to the justificatory question reveals what those reasons are: This person participates in a legitimate social practice which is

structured by certain norms, one of which is EN. Thus, they have good reason to accept EN.

We can conclude from this that EN is a correct epistemic norm of belief. Notice that simply *accepting* EN will not prevent one from being subjected to the responses captured in **Epistemic Accountability** (not only because there are other norms which one should also accept – more on this momentarily). There might be instances where a person accepts EN but they are non-culpably mistaken about the objective evidential situation vis-à-vis some topic or issue. For instance, this person might (blamelessly) *think* that some proposition constitutes strong evidence E for the truth of p, when in fact it doesn't. They might then treat that proposition as a warranting reason to believe p. In such cases, we might still subject the person to the responses captured in **Epistemic Accountability**. In other words, acceptance of the relevant norms isn't *enough* to stave off the responses captured in **Epistemic Accountability**. One also needs objectively accurate views.

This brings us to the following question: What *other* norms structure our practice of epistemic accountability? In Chapter 2 I noted that EN is somewhat ecumenical as far as its content goes: It incorporates an externalist view of "strong evidence", but it also incorporates various internalist elements such as "possession" and a requirement that the evidential connection be "clear" to the subject. How does my approach relate to the dispute between internalists and externalists when it comes to the content of epistemic norms? As I've already indicated above, my approach will incorporate many externalist elements. We saw this above when considering the case of Bill in **Case 2**. In maintaining false views about climate change, Bill isn't manifesting a bad "doxastic self"; we could even say that he's *accepted* various externalist norms of belief which require truth or knowledge. Nevertheless, I have argued that it

is appropriate to take up the responses captured in **Epistemic Accountability** towards Bill in a way that goes beyond mere assessment vis-à-vis a standard. Given the first condition of **Epistemic Accountability** ("(i) A judges that B violates a norm of the relevant kind") this means that Bill will have violated norm.

What norm has Bill violated in believing what he does? In answering this question, we will be attempting to fill out the content of the norms which structure our practice of epistemic accountability. As I noted in the previous chapter, we can think of the forms of response captured in **Epistemic Accountability** as ways of responding to "claim-makers". Thus, the norms which structure our practice of epistemic accountability can be thought of as delivering criteria for *successful* "claim-makers". What, then, are the norm(s) that Bill runs afoul of which render him an *unsuccessful* claim-maker when it comes to the topic of climate change? Recall that the fact that Bill is unsuccessful in this regard is revealed in the fact that we can appropriately take up the responses captured in **Epistemic Accountability** towards him. A natural suggestion is that Bill flout's a *truth* norm. In other words, his beliefs regarding climate change run afoul of norms that require our beliefs to be true. Consider, for instance, TN:

TN: Don't believe p unless p is true.

Perhaps Bill is an unsuccessful claim-maker vis-à-vis climate change since his beliefs regarding this subject matter aren't true. The thought here is that, when we make "claims" regarding some topic or subject matter, we *express our beliefs*[177]. Thus, in order for our claims to be fully

[177] But what about instances where someone claims p but *isn't* being sincere? Won't they be an unsuccessful claim-maker? And if so, are prohibitions on lying "epistemic" norms according to my view? Prohibitions on lying are not epistemic norms according to my view. This goes back to issues discussed in Chapter 4. Specifically, it doesn't seem that the way that we respond to liars

successful, the beliefs which are thereby expressed must be true.

This suggestion is on the right track, but it actually doesn't go far enough. In other words, it seems that, in order to be a fully successful claim-maker, the beliefs which are thereby expressed have to meet *more* than just a truth norm. To see this, let's return to the example of TOM AT THE MOVIES. Recall that Tom was the person who also culpably denies the existence of climate change. Here's the relevant case:

> TOM AT THE MOVIES: Tom saw a movie at the downtown cinema this past Tuesday at 7:30 p.m. On the basis of this, Tom believes that the same movie will be playing at the same time at the downtown cinema the following Tuesday. In other words, without checking to see if the movie is even still playing at the downtown cinema the following week, or whether it's still playing but the screening days/times have changed, etc., Tom comes to believe, just on the basis of his seeing the movie on Tuesday at 7:30 p.m., that the same movie will be screened again the following Tuesday at the same time.

Let's say that, as it turns out, Tom's belief is *true*; the movie that he saw will indeed by screened again the following Tuesday at the same time. Let's also say that there's an abundance of evidence for the truth of this claim: the schedule is posted on the theater's website, the theater employees can attest to the fact that the movie will be screened on that day/time, etc. However, Tom doesn't base his belief on any of this. Rather, he bases his belief on nothing more than the fact that he saw the movie on that particular Tuesday at 7:30 p.m.

Even though Tom's belief *doesn't* run afoul of TN, it doesn't seem that he is thereby a successful claim-maker when it comes to the screening days/times for the relevant movie. Imagine, for instance, that we're curious whether or not the movie that Tom saw is playing again at the downtown cinema next Tuesday at 9:30 and we overhear Tom telling someone that it is.

---

is *distinctly epistemic*. Rather, it seems that holding an individual *morally* accountable is the appropriate form of response in such instances.

Say that this person asks Tom: "How do you know that?", and Tom adduces his meager support. As a result, we can lower our trust in Tom when it comes to this subject matter. However, imagine that we look up the screening times and we discover that Tom was right after all. The appropriate response here is *not* the following: "We should have trusted him!". Even after finding out that Tom's belief was true, we still recognize that we were fully justified in responding to him as we did. Moreover, even *after* we find out that Tom's beliefs are true, we will *still* not take his claims at face value when it comes to the relevant subject matter, so long as he continues to base them on the meager support mentioned above. In other words, even after the fact, we can still take up the responses captured in **Epistemic Accountability** vis-à-vis Tom when it comes to his belief about the screening days/times for the movie. Even though we know that Tom's *belief that p* is true, we no longer see *Tom's claims that p* as good reason to think that the proposition which is *presented* as true actually *is* true. Recall that this was one central way in which the modifications of trust captured in **Epistemic Accountability** can be manifested; we can cease to take a person's claims at face value, where this involves ceasing to treat the fact that the person *presented* some proposition as true as good reason to think that the proposition actually *is* true. In the case of Tom, we have *other* grounds for being confident in the truth of the relevant proposition; grounds that are unrelated to Tom's claiming that it's true. The case at hand relates to an important distinction between believing *a person* and believing *a proposition*. Consider, for instance, Moran (2018):

> I want to examine the relation of believing where its direct object is not a proposition but a person. For in the basic case…it is the speaker who is believed, and belief in the proposition asserted follows from this. These are different epistemic phenomena. For the hearer might not believe the speaker at all, taking her for a con artist, but yet believe that

> what she has said is in fact true. Whereas when the hearer believes the speaker, he not only believes what is said but does so on the basis of taking the speaker's word for it.[178]

Moran's distinction between believing a *person* and believing a *proposition* allows us to understand our responses to Tom above. Even after we find out that Tom's beliefs are true, we still don't believe *him*, i.e. we don't take his claims as good reason to think that the proposition which is presented as true actually is true. This is so even though we believe that *what he says* (i.e. the proposition expressed) is true.

The cases of Bill and Tom show us that the truth of one's expressed beliefs is a *necessary but not sufficient* condition for being a successful claim-maker. Once again, the case of Tom shows us that, even when a person's beliefs regarding some topic or subject matter are *true* and *we know them to be true*, we can *still* modify our trust and reliance in them in the ways captured in **Epistemic Accountability**. What, then, does a *successful* claim-maker look like? The cases of Bill and Tom are suggestive of the following: We don't *just* want claim-makers who have accurate or true beliefs (Tom meets this criterion), and we don't *just* want claim-makers with good "doxastic selves" (Bill meets this criterion). Rather, it seems that what we want are claim-makers who have true or accurate beliefs *because* they have good "doxastic selves". Importantly, it doesn't seem that having true beliefs that *express* or *disclose* a good doxastic self will be enough. Rather, we want the person to have their true beliefs *because* of their good doxastic self. Consider, for instance, a classic "Gettier case" from Roderick Chisholm:

> SHEEP: Charles is standing next to a field. In the middle of the field, he sees what looks exactly like a sheep. Charles comes to believe that there is a sheep in the field. In fact, the "sheep" is not a sheep at all but a very convincingly disguised dog. However, there is a

[178] Moran (2018) p. 38.

> sheep in the field; it is behind a hill in the middle of the field where Charles cannot see it and he has no evidence for it.[179]

Focus on Charles' belief that there's a sheep in the field. Charles' belief is true. Moreover, let's also say that Charles, not unlike Bill, has an impeccable "doxastic self"; he tries his best to conform his beliefs to the evidence, he makes a concerted effort to avoid various biases, and he is diligent and scrupulous when taking on new beliefs. In spite of all of this, something has gone awry in SHEEP. It's not the case the Charles' (true) belief fails to *express* or *disclose* a good doxastic self; his belief does precisely that. Rather, the issue is that Charles' true belief isn't connected in the right way with his doxastic self; he doesn't have the true belief *because* of his doxastic self. Rather, he has the true belief simply as a result of luck. Once again, even knowing all of the details of the case, we can still subject Charles to the forms of response captured in **Epistemic Accountability**. In other words, we can cease to take Charles' claims that there's a sheep in the field at face value. We can do this even though Charles' expressed belief is true, *and we know it to be true*.

Perhaps, then, when it comes to claim-makers, what we want are individuals who have true beliefs *not* as a result of luck. One way to eliminate luck would be to require that our true beliefs *also* be formed by belief-forming processing that are *de facto* reliable in the environments in which we find ourselves. However, once again, if we stick to our responses in **Epistemic Accountability**, it does not seem that this will work. Consider one final stock example from the epistemology literature, Laurence BonJour's case of Norman the clairvoyant:

> CLAIRVOYANCE: Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He

[179] Chisholm (1977).

> possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.[180]

Imagine that we know about Norman's clairvoyant powers. Even if this is the case, we can *still* subject him to the responses captured in **Epistemic Accountability** when it comes to his belief about the President's whereabouts. We wouldn't take *Norman's claim* that the President is in New York City as a good reason to believe the truth of the relevant proposition. Rather, we have other grounds, unrelated to the fact that *Norman claimed such-and-so,* which underwrite our confidence in the truth of the proposition. Even though we believe *the proposition*, we don't believe *the person*.

Where does this leave us when it comes to the norms which structure our practice of epistemic accountability? What, in other words, are the relevant criteria for successful claim-makers? What we seem to want in a claim-maker is someone who has true beliefs *because* of their good doxastic self, where the latter has *something* to do with the elimination of luck, but where this elimination *isn't* simply a matter of having true beliefs that are caused by processes that are *de facto* reliable in the environments in which we find ourselves. The concept of 'knowledge' seems to be the best candidate to characterize beliefs of that satisfy these desiderata. If this is right, then it seems that our practice of epistemic accountability will be structured by *knowledge* norms, e.g. KN:

KN: Don't believe p unless you know p.

This is perhaps the norm that is flouted in the cases of SHEEP and CLAIRVOYANCE, and it is

[180] BonJour (1980) p. 62.

perhaps also the norm which underwrites our modification of trust in the individuals in these cases. If this suggestion is on the right track, then looking more closely at the different situations in which taking up the responses captured in **Epistemic Accountability** towards a person are appropriate might help us illuminate the conditions that have to be met in order to count as knowing something. Additionally, as I've indicated, the above suggests that our practice of epistemic accountability is structured by knowledge norms such as KN. If this is correct, then participants in the social practice of epistemic accountability will have good reason to accept knowledge norms, i.e. they will have good reason to accept norms which require their beliefs to amount to knowledge.

Conclusion

I began this dissertation by considering a view that I labeled "evidential minimalism". According to evidential minimalists, evidence *itself* is normative for belief. *Non*-minimalists, by contrast, hold that something more substantial needs to be on the table, over and above strong evidence for the truth of p and certain other minimal conditions, in order for there to be a normative reason for belief. I argued against the minimalist position. However, I *also* argued against a number of non-minimalist views that are popular in the contemporary literature on epistemic normativity, e.g. Jane Friedman's "zetetic" approach and various versions of epistemic instrumentalism.

What emerged from the discussion is a novel and attractive non-minimalist position regarding evidential normativity. According to my view, the normativity of evidence is best understood internal to a social practice of a particular kind, viz. our practice of epistemic accountability. This position is non-minimalist insofar as it locates the normativity of evidence not in evidence itself, or in some other "purely cognitive" consideration, but rather in a social practice of a certain kind. Nevertheless, the approach developed here also has the resources to vindicate a very wide scope of application for epistemic norms, in addition to providing a basis for their authoritative strength. I take these to be virtues of the account on offer here that other forms of non-minimalism have difficulty attaining.

In addition to providing us with an understanding of the normativity of evidence, the account developed here also vindicates a broadly *externalist* approach to epistemic norms of belief in general. Thus, in spite of the fact that the account on offer here bases epistemic norms in *social practices* – things which may seem unwieldy, excessively varied, or inherently flawed – I

I have vindicated a robustly externalist approach to epistemic norms; one which privileges not just truth but *knowledge*. Finally, the account on offer provides us with a novel way of developing an answer to the "justificatory" question vis-à-vis our acceptance of epistemic norms of belief. As I have indicated, the "justificatory" question is a first-personal question regarding the reasons which ground one's acceptance of a norm or class of norms. The answer that I provide does not appeal to some consideration which is independent of our actual social practices and the psychological facts which render us "participants" within those practices. Rather, according to my view, the reason why we should accept epistemic norms is simply that they are binding on us. These norms are binding on us, once again, given that they structure a legitimate social practice that we participate in.

According to this approach, it would be apt to say that we "find ourselves" as participants within certain practices. This is not to say that we are hostage to our existing social world. It is simply to say that this is our unavoidable starting position. Adjustments to our practices are made from the inside out, in a piecemeal and incremental fashion. Where there is no reason to reject an existing practice in its entirety, it is thereby legitimate and its participants are bound by the norms and expectations which structure its activities.

Bibliography

Abramson, K., Leite, A. (2017). "Love, Value, and Reasons" In C. Grau, A. Smuts (Eds.) *The Oxford Handbook of Philosophy of Love.* New York: Oxford University Press.

Alston, W. (1988) "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2:257-299 (1988).

Basu, R. (2019a). "The Wrongs of Racist Beliefs", *Philosophical Studies* 176(9):2497–2515.

------. (2019b). "What We Epistemically Owe To Each Other", *Philosophical Studies* 176(4):915-931.

Berker, S. (2013). "Epistemic Teleology and the Separateness of Propositions". *Philosophical Review,* 122, 337-93.

Bonjour, L. (1980). "Externalist Theories of Empirical Knowledge*." Midwest Studies in Philosophy* 5 (1):53-73.

Boult, C. (2021a). "Epistemic Blame". *Philosophy Compass.*

------. (2021b). "Standing to Epistemically Blame". *Synthese.* https://doi-org.proxyiub.uits.iu.edu/10.1007/s11229-021-03293-3

------. (2021c). "The Significance of Epistemic Blame". *Erkenntnis.* https://doi-org.proxyiub.uits.iu.edu/10.1007/s10670-021-00382-0

------. (2020). "There is a Distinctly Epistemic Kind of Blame". *Philosophy and Phenomenological Research.* https://doi-org.proxyiub.uits.iu.edu/10.1111/phpr.12726

------. (2017). "Epistemic Normativity and the Justification-Excuse Distinction". *Synthese* 194, 4065-4081.

Brandt, R. (1979). *A Theory of the Good and the Right.* New York: Oxford University Press.

Brown, J. (2018). "What is Epistemic Blame?" *Nous.* https://doi.org/10.1111/nous.12270

------. (2020). "Epistemically Blameworthy Belief". *Philosophical Studies* 177 (12):3595-3614.

Buckley, D. (2021). "Varieties of Epistemic Instrumentalism". *Synthese.* 198, 9293-9313.

------. (Forthcoming) "Against Evidential Minimalism". *Episteme.*

Chan, T. (2014) Eds. *The Aim of Belief.* Oxford University Press.

Chisholm, R. (1977). *Theory of Knowledge*. Prentice Hall.

Cohen, S. (1984). "Justification and Truth". *Philosophical Studies*. 46:279-96.

------. (2016). "Theorizing about the Epistemic". *Inquiry*, 59, 839-857.

Côté-Bouchard, C. (2017). "Epistemological Closed Questions: A Reply to Greco". *Manuscrito*, 97-111.

------. (2015). "Epistemic Instrumentalism and the Too Few Reasons Objection." *International Journal of Philosophical Studies,* 23(3), 337-355.

------. (2016). "Can the Aim of Belief Ground Epistemic Normativity?" *Philosophical Studies,* 173(12), 3181-3198.

Cowie, C. (2014) "In Defense of Instrumentalism about Epistemic Normativity." *Synthese,* 191(16), 4003-4017.

Craig, E. (1990). Knowledge and the State of Nature. Oxford University Press.

Dyke, M. (2020). "Could Our Epistemic Reasons be Collective Practical Reasons?". Nous, 55, 842-862.

Enoch, D. (2006). "Agency, Shmagency: Why Normativity won't come from what's Constitutive of Action.", *The Philosophical Review*, 115, 169-198.

------. (2011). "Shmagency Revisited" in M. Brady (ed.), *New Waves in Metaethics.* London: Palgrave-Macmillan.

Faulkner, P. (2011). *Knowledge on Trust*. New York: Oxford University Press.

Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press.

Frankurt, H. (2003). "Freedom of the Will and the Concept of a Person" in G. Watson (ed.) *Free Will.* Oxford University Press.

Fratantonio, G. (Forthcoming). "Evidential Internalism and Evidential Externalism" in M. Lasonen-Aarnio and C. Littlejohn (eds.), *The Routledge Handbook for the Philosophy of Evidence.*

Friedman, J. (2020). “The Epistemic and the Zetetic”. *The Philosophical Review*.

------. (2018). “Inquiry and Beliefs.” Nous.

------. (Forthcoming). “Zetetic Epistemology”. *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere.* Baron Reed & A. K. Flowerree (eds.).

Fricker, E. (2006). “Second Hand Knowledge”. *Philosophy and Phenomenological Research*, 73, 592-618.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Greco, D. (2019). “Justification and Excuses in Epistemology”. *Nous.* 1-21.

------. (2015). “Epistemological Open Questions”. *Australasian Journal of Philosophy*, 93, 509-523.

Heathwood, C. (2009). “Moral and Epistemic Open-Question Arguments”. *Philosophical Books*, 50, 83-98.

Hieronymi, P. (2008). “Responsibility for Believing”. *Synthese,* 161, no. 3: 357-373.

Hinchman, E. (2005). “Telling as Inviting to Trust” *Philosophy and Phenomenological Research*. 70 (3): 562-87.

------. (2014). “Assurance and Warrant.” *Philosophers Imprint.* 14(17).

Hofmann, F. (2021). “Is Evidence Normative?” *Philosophia*. 49, 667-684.

Horst, D. (Forthcoming). “In Defense of Constitutivism about Epistemic Normativity”. *Pacific Philosophical Quarterly.*

Kauppinen, A. (2018). “Epistemic Norms and Epistemic Accountability” *Philosophers Imprint.* Vol. 18, no. 8.

Kelly, T. (2003). “Epistemic Rationality as Instrumental Rationality: A Critique.” *Philosophy and Phenomenological Research*, 66(3), 612-640.

------. (2007). “Evidence and Normativity: Reply to Leite.” *Philosophy and Phenomenological Research*, 75(2), 465-474.

------. (2014). “Evidence”. *Stanford Encyclopedia of Philosophy*

Kiesewetter, B. (2021). "Are Epistemic Reasons Normative?". *Nous*, 1-26. https://doi.org/10.1111/nous.12379.

------. (2017). *The Normativity of Rationality*. Oxford University Press.

Kornblith, H. (2002). *Knowledge and its Place in Nature*. New York: Oxford University Press.

------. (2022). "Against Strawsonian Epistemology: Testimony, Self-Knowledge, Promising, and Resolving," in N. Ballantyne and D. Dunning, eds., *Reason, Bias, and Inquiry: The Crossroads of Epistemology and Psychology*, Oxford University Press, 254-275.

Korsgaard, C. (1996) *The Sources of Normativity*. Cambridge University Press.

------. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press.

Lacey, N & Pickard, H. (2021). "Why Standing to Blame May be Lost but Authority to Hold Accountable Retained." *The Monist*. 104, 265-280.

Lackey, J. (2008). *Learning from Words*. New York: Oxford University Press.

Laudan, L. (1990). "Normative Naturalism". *Philosophy of Science*, 57, 44-59.

Leite, A. (2013). "But That's Not Evidence; It's Not Even True!". *Philosophical Quarterly* 63 (250): 81-104.

------. (2007). "Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research*. 75(2), 456-464.

Littlejohn, C. (2013). "No Evidence is False". *Acta Analytica*, 28: 145-159.

------. (Forthcoming). "A Plea for Epistemic Excuses," in F. Dorsch and J. Dutant (eds.) *The New Evil Demon*. Oxford University Press.

Lockard, M. (2013). "Epistemic Instrumentalism." *Synthese*, 190(9), 1701-1718.

Maguire, B. & Woods, J. (2020). "The Game of Belief". *The Philosophical Review*, 129, 211-49.

Manne, K. (2016). "Democratizing Humeanism" In E. Lord, B. Maguire (Eds.) *Weighing Reasons* (pp. 123-140). New York: Oxford University Press.

Mantel, S. (2019). "Do Epistemic Reasons Bear on the Ought Simpliciter?" *Philosophical Issue*, 29, 214-

27.

Marušić & White (2018). "How Can Beliefs Wrong" A Strawsonian Epistemology". *Philosophical Topics*, 46, 97-114.

McCormick, M.S. (2020). "Believing Badly: Doxastic Duties are Not Epistemic Duties". In K. McKain & S. Stapleford (Eds.), *Epistemic Duties: New Arguments, New Angles* (pp. 29-43). Routledge.

McHugh (2012). "Epistemic Deontology and Voluntariness" *Erkenntnis*, 77, 65-94.

McMyler, B. (2011). *Testimony, Trust, and Authority*. New York: Oxford University Press.

Moore, G.E. (1903). *Principia Ethica*. Cambridge University Press.

Moran, R. (2006) "Getting Told and Being Believed" in The Epistemology of Testimony, ed. J. Lackey and E. Sosa. Oxford: Clarendon Press.

------. (2018). *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oxford University Press.

Nelson, M. (2010). "We Have no Positive Epistemic Duties." *Mind*, vol. 119, 473, 83-102.

Nottelmann, N. (2007). *Blameworthy Belief: A Study in Epistemic Deontologism*. Springer.

Paakkunainen, H. (2018). "Doubts about "Genuinely Normative" Epistemic Reasons." In C. McHugh, J. Way, D. Whiting (Eds.) *Metaepistemology* (pp. 122-140). New York: Oxford University Press.

Papineau, D. (2013). "There are no Norms of Belief". In T. Chan (ed.), *The Aim of Belief*. Oxford University Press.

------. (2003). *The Roots of Reason: Philosophical Essays on Rationality, Evolution, and Probability*. Oxford University Press.

Parfit, D. (2011). *On What Matters*. Vol. 1. *Oxford University Press*.

Piovarchy, A. (2021). "What do We Want from a Theory of Epistemic Blame?", *Australasian Journal of Philosophy,* Vol. 99, no. 4, 791-805.

Quine, W.V. (1986). "Reply to Morton White" in Lewis Hahn and Paul Schilpp (eds.), *The Philosophy of W.V. Quine,* 663-5.

Reis-Dennis, S. (2018). "Responsibility and the Shallow Self", *Philosophical Studies,* 175, 483-501.

Rettler, L. (2018). "In Defense of Doxastic Blame". *Synthese*, 195: 2205-2226.

Rinard, S. (2015). "Against the New Evidentialists" *Philosophical Issue*, 25, 208-23.

Ross, A. (1986). "Why do we Believe what we are Told? *Ratio* (1): 69-88.

Shah. N. (2003). "How Truth Governs Beliefs" *The Philosophical Review*. 112, 447-82.

Shah, N. and Velleman, J.D. (2005). "Doxastic Deliberation". *The Philosophical Review*, 114, 497-534.

Scanlon, T.M. (1998). *What we Owe to Each Other*. Harvard University Press.

------. (2013). "Interpreting Blame" in D.J. Coates & N.A. Tognazzini (eds.), *Blame: Its Nature and Norms*. Oxford University Press.

------. (2008). *Moral dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press.

Schmidt, S. (2021). "Epistemic Blame and the Normativity of Evidence". *Erkenntnis*. https://doi-org.proxyiub.uits.iu.edu/10.1007/s10670-021-00430-9

Schmitt, F. (2010). "The Assurance View of Testimony". in A. Haddock, A. Millar, and D. Pritchard, eds., *Social Epistemology*. Oxford University Press.

Schwitzgebel, E. (2011). *Perplexities of Consciousness*. MIT press.

------. (2008). "The Unreliability of Naïve Introspection". *Philosophical Review*, 117, 245-273.

Sharadin, N. (2018). "Epistemic Instrumentalism and the Reason to Believe in Accordance with the Evidence." *Synthese*,195(9), 3791-3809.

Sher, G. (2006). *In Praise of Blame*. Oxford University Press.

Shoemaker, D. (2015). Ecumenical attributability. In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of moral responsibility: New essays*. New York: Oxford University Press.

Silins, N. (2005). "Deception and Evidence". *Philosophical Perspectives*, 19, 375-404.

Smith, A. (2012). "Attributability, Answerability, and Accountability: In Defense of a Unified Account," *Ethics* 122, pp. 575–589.

------. (2015) “Responsibility as Answerability,” *Inquiry* 58:2: 99-126.

------. (2019). "Who's Afraid of a Little Resentment?" in *Oxford Studies in Agency and Responsibility,* vol. 6.

------. (2008). “Control, Responsibility, and Moral Assessment,” *Philosophical Studies* 138:3 367-392.

------. (2005). “Responsibility for Attitudes: Activity and Passivity in Mental Life,” *Ethics* 115, pp. 236–271.

Srinivasan, A. (2015). “Normativity Without Cartesian Privilege”. *Philosophical Issues*. https://doi.org/10.1111/phis.12059

Steglich-Petersen, A. (2018). “Epistemic Instrumentalism, Permissibility, and Reasons for Belief.” In C. McHugh, J. Way, D. Whiting (Eds.) *Normativity: Epistemic and Practical* (pp. 260-280). New York: Oxford University Press.

------. (2006). “No Norm Needed: On the Aim of Belief”. *Philosophical Quarterly.* 56, 499-516.

------. (2009). “Weighing the Aim of Belief”. *Philosophical Studies.* 145, 395-405.

Strawson, P. (1962/2003). “Freedom and Resentment”. Proceedings of the British Academy 48, 1–25. Reprinted in Gary Watson (ed.) Free Will. Oxford: Oxford University Press, 72–93.

Sripada, C. (2016). “Self-Expression: A Deep-Self Theory of Moral Responsibility”, *Philosophical Studies,* 173, 1203-1232.

Taylor, S., Brown, J. (1988) “Illusion and Well-Being: A Social Psychological Perspective on Mental Health” *Psychological Bulletin,* 103(2), 193-210.

Velleman, D. (2000). “On the Aim of Belief” in his *The possibility of Practical Reason*. Oxford University Press.

Wallace, R.J. (2011). “Dispassionate Opprobrium”. In R.J Wallace, R. Kuhmar, & S. Freeman (Eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: OUP.

Watson, G. (2004). "Two Faces of Responsibility" in *Agency and Answerability*. Oxford University Press.

------. (2003). "Free Agency" in G. Watson (ed.) *Free Will.* Oxford University Press.

Wedgwood, R. (2002). "The Aim of Belief". *Philosophical Perspectives,* 16, 267-97.

Whiting, D. (2013). "Truth: The Aim and Norm of Belief". *Teorema,* vol. 32, 3, 121-135.

Williams, B. (1973). *Problems of the Self.* Cambridge University Press.

Williamson, T. (2000). *Knowledge and its Limits.* Oxford University Press.

------. (Forthcoming). "Justification, Excuses, and Skeptical Scenarios". In F. Dorsch and J. Dutant (eds.) *The New Evil Demon*. Oxford University Press.

Wolf, S. (2011). "Blame, Italian Style". In R.J Wallace, R. Kuhmar, & S. Freeman (Eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford: OUP.

------. (2019). "Attributability and the Self". In *Oxford Studies in Agency and Responsibility,* Vol. 5.

------. (1990) *Freedom within Reason*. Oxford University Press.

Daniel S. Buckley
1635 Fairchild Avenue, Manhattan, KS 66502
Dbuckley86@gmail.com
(484) 686-8323

Education

Ph.D. Philosophy, Indiana University, Bloomington, 2022
Minor in History and Philosophy of Science

M.A. Philosophy, Northern Illinois University, 2013

B.A. Political Science, *magna cum laude*, Elizabethtown College, 2008
Minor in Asian Studies
Academic Year 2007-2008 study abroad in Beijing, China

Areas of Specialization and Competence

SPECIALIZATION: Epistemology

COMPETENCE: Ethics, Metaethics, Applied Ethics, History of Philosophy (esp. Modern, 20th Century Analytic, and Existentialism), Social and Political Philosophy, Philosophy of Mind

Academic Employment

Visiting Assistant Professor, Kansas State University, 2022 – present
Future Faculty Teaching Fellow, Butler University, 2021-2022

Research

PUBLICATIONS

1. Buckley, D. (Forthcoming) "Against Evidential Minimalism". *Episteme*
2. Buckley, D. (2021) "Varieties of Epistemic Instrumentalism". *Synthese,* 198:9293-9313. DOI: 10.1007/s11229-020-02634-y. Special Issue on Instrumentalism about Epistemic Rationality.

WORKS IN PROGRESS

"Epistemic Accountability without Epistemic Blame" (in preparation)
"Evidence and Epistemic Normativity: A Social-Accountability Approach" (in preparation)
"Does Epistemology Have a Normative Question?" (in preparation)

PRESENTATIONS

"A Social-Accountability Approach to Epistemic Normativity"
- York College of Pennsylvania, March 2022 (invited)

"Epistemic Accountability and Externalist Norms of Belief"
- American Philosophical Association, Central Division Meeting. Chicago, IL February

2022 (peer-reviewed)
- Indiana Philosophical Association, Fall Meeting. Bloomington, IN. November 2021. (peer-reviewed)

"Epistemic Norms and Holding Ourselves (and Each Other) Epistemically Accountable"
- September 18th, 2020. Nelson Symposium Lecture. Indiana University, Bloomington, IN (online)

Commentary on Michael Milona's "What Cases of Appropriation (Probably) Reveal about the Content and Taboo Nature of Unappropriated Slurs"
- 6th Annual Northern Graduate Philosophy Conference, Northern Illinois University. April 19-20, 2013

Dissertation

Title: Evidence and Epistemic Normativity

Committee: Adam Leite (chair), Kate Abramson, Jordi Cat, and Gary Ebbs

Short Abstract: Evidence is often taken to be "normative" for doxastic agents. What accounts for the normativity of evidence? My dissertation is devoted to answering this question. In order to answer it, I suggest we look to the ways in which we relate to *others* in their capacity as epistemic agents. I argue that, when we attend to our actual social-epistemic practices, we will see that we hold each other *accountable* in a distinctly epistemic way. I argue that our practice of epistemic accountability is a legitimate social practice and that it is structured by certain evidential norms of belief. Evidential norms have authority over us, then, given our participation in a social practice that includes a distinctly epistemic form of interpersonal accountability.

Honors and Awards

APA Graduate Student Travel Stipend, 2022
Future Faculty Teaching Fellowship, 2021-2022
James B. Nelson Dissertation Year Fellowship, 2019-2020
Outstanding Associate Instructor Award, Indiana University, Department of Philosophy, Spring 2019
Bo and Lynn Clark Essay Prize, Indiana University, Department of Philosophy, 2018
Bo and Lynn Clark Essay Prize (for first or second year graduate students), Indiana University Department of Philosophy, 2014
College Scholar, Elizabethtown College (awarded to students with a cumulative GPA of 3.75 or higher after completing 60 or more credits)
Pi Sigma Alpha, Elizabethtown College
Dean's Scholarship, Elizabethtown College

Teaching

COURSES TAUGHT AS SOLE INSTRUCTOR

Kansas State University
Biomedical Ethics (multiple sections)

Butler University

First-Year Seminar: The Things we Believe
Introduction to Symbolic Logic
History of Modern Philosophy

Indiana University, Bloomington

Introduction to Ethics (4x)
Applied Ethics (Intensive Writing Course)
Introduction to Social and Political Philosophy

COURSES TAUGHT AS TEACHING ASSISTANT

Indiana University, Bloomington

Introduction to Ethics (Marcia Baron)
Introduction to Philosophy (Adam Leite)
Introduction to Existentialism (Allen Wood)
Elementary Logic (Mark Kaplan)
Critical Thinking (2x; Leah Savion)
Public Oral Communications (John Arthos; online)

Northern Illinois University

Biomedical Ethics (Sharon Sytsma)
Introduction to Metaphysics (Alicia Finch)
Logic and Critical Reasoning (2x; John Beaudoin)
Contemporary Moral Issues (2x; Nicoleta Apostol)

GRADER

Indiana University, Bloomington

Introduction to Symbolic Logic (Joan Weiner)
Introduction to the Philosophy of Mind (Frederick Schmitt)

GUEST LECTURES

York College of Pennsylvania

"Relevance", Spring 2022, Critical Thinking (Dennis Weiss)

Indiana University, Bloomington

"Moritz Schlick's *Positivism and Realism*" Fall 2015, Introduction to Philosophy (Adam Leite)
"Quine's Two Dogmas of Empiricism" Spring 2015, graduate seminar in philosophy of science (Jordi Cat)

Research Assistantship

For Fredrick Schmitt (Fall 2020). Reviewed and summarized recent literature in the area of social epistemology.

Non-Academic Employment

Paralegal, 2008-2009
Marshall, Dennehy, Warner, Coleman, and Goggin. King of Prussia, PA
Worked directly with attorneys representing defendants in medical malpractice and professional liability claims.

Academic Service

Referee, *Philosophical Quarterly* (x1)
Faculty Liaison for the Undergraduate Philosophy Club at Butler University, 2021-22.
Graduate Student Colloquium Coordinator, IUB philosophy, 2018-2019.
IUB Dept. of Philosophy Representative to the Graduate and Professional Student Organization (GPSO), 2014-2015.
Member of IUB Graduate and Professional Student Organization (GPSO) Committee on Student Affairs, 2014-2015.
Referee, 6th Annual Northern Graduate Philosophy Conference, Northern Illinois University.

Workshops and Pedagogical Training

Teaching for Student Success: An Evidence-Based Approach, Indiana University, Bloomington, IN. Four-week pedagogy course covering course design, methods of assessment, and effective teaching practices. June 2021 (online).

American Association of Philosophy Teachers (AAPT) one-day workshop on inclusive pedagogy. September 2017. Bloomington, IN.

Training to teach a writing-intensive summer course (Applied Ethics) at Indiana University's Center for Innovative Teaching and Learning (CITL).

Graduate Coursework (*=Audit †= Undergraduate Course)

HISTORY OF PHILOSOPHY

Plato on Knowledge (Pieter Sjoerd Hasper, Indiana)
Kant's *Critique of Pure Reason* (Allen Wood, Indiana)
Hume's *A Treatise of Human Nature* (Frederick Schmitt, Indiana)
Frege (Joan Weiner, Indiana)
Carnap and Quine (Gary Ebbs, Indiana)
History of 20th Century Analytic Philosophy (David Buller, Northern Illinois)
*†Wittgenstein (Kirk Ludwig, Indiana)

METAPHYSICS AND EPISTEMOLOGY

The Internal and External in Epistemology (Mark Kaplan, Indiana)
External World Skepticism (Adam Leite, Indiana)
Collective Action (Kirk Ludwig, Indiana)
Epistemology of Perception (Geoff Pynn, Northern Illinois)
Metametaphysics (Alicia Finch, Northern Illinois)
*Wittgenstein's *On Certainty* and Contemporary Epistemology (Adam Leite, Indiana)
*Self-Knowledge (Adam Leite, Indiana)

*Contemporary Epistemology (Geoff Pynn, Northern Illinois)
*†Knowledge, Skepticism, and Context Sensitivity (Lenny Clapp, Northern Illinois)

LOGIC AND THE PHILOSOPHY OF LANGUAGE

Graduate Logic I (David McCarty, Indiana)
Intermediate Logic (David Buller, Northern Illinois)
Philosophy of Language (Tomis Kapitan, Northern Illinois)
Truth (David Buller, Northern Illinois)
Truth Conditional Pragmatics (Lenny Clapp, Northern Illinois)
*Truth by Linguistic Convention, Analyticity, Apriority, and Contextual Apriority (Gary Ebbs, Indiana)

PHILOSOPHY OF SCIENCE AND MIND

History and Philosophy of Comparative Cognition (Colin Allen, Indiana)
Pragmatism and Philosophy of Science (Elisabeth Lloyd, Indiana)
Modern Philosophy of Science (Jordi Cat, Indiana)
Philosophy of Mind (Carl Gillett, Northern Illinois)
Metaphysics of Science (Carl Gillett, Northern Illinois)

ETHICS AND THE HISTORY OF ETHICS

Contemporary Ethical Theories (Marcia Baron, Indiana)
Aristotle's Ethics and Politics (Pieter Sjoerd Hasper, Indiana)
Advanced Ethical Theory (Jason Hanna, Northern Illinois)
Kant's Ethics (Sharon Sytsma, Northern Illinois)
*Legal Philosophy (Marcia Baron, Indiana)
*Topics in Contemporary Ethics: Blame (Kate Abramson, Indiana)
*Human Psychology and the Claims of Ethics (Kate Abramson, Indiana)
*Kant's Ethics (Allen Wood, Indiana)

.